The Full, Persistent, and Symmetric Pass-Through of a Temporary VAT Cut*

Tiago Bernardino†
IIES, Stockholm University

Ricardo Duque Gabriel‡
Federal Reserve Board

João Quelhas§
Stockholm University

Márcia Silva-Pereira¶
Nova SBE

May 2025

Abstract

We investigate the pass-through of a temporary value-added tax (VAT) cut on selected food products to consumer prices in Portugal. Exploiting a novel data set of daily online prices, we find that the VAT cut was fully transmitted to consumer prices, persisted throughout the policy duration, and prices returned to the pre-implementation trend after reversal. We discuss two potential mechanisms driving this result: the policy's salience to consumers in a high-inflation environment and the decline of producer prices when implemented. We estimate that the policy reduced the inflation rate by 0.68 percentage points on impact.

JEL classification: E31, H20, H22

Keywords: value-added tax; pass-through; microdata; public policy; inflation

*We thank the editor, Antoine Bozio, two anonymous referees, Youssef Benzarti, Daniel Dias, Mitch Downey, Rita de la Feria, Erwan Gautier, Andrew Goodman-Bacon, Jarkko Harju, Camille Landais, Kieran Larkin, Jens Oehlen, Elvira Prades, Stephan Schneider, David Schönholzer, David Seim, Jeffrey Wooldridge, and participants at the 17^{th} PEJ Annual Meeting, the 12^{th} LuBraMacro Meeting, the 8^{th} Swedish Conference in Economics, the 3^{rd} "Portuguese Around the World: Central Banking Edition" Conference, and seminars at Banco de Portugal, ECB, FED Board, GEE/GPEARI, IIES, ISEG, Nova SBE, U. Minho, U. Porto for valuable comments and discussions. We thank Banco de Portugal's Microdata Research Laboratory for their data support. T. Bernardino acknowledges financial support from the Jan Wallanders & Tom Hedelius Foundation (H21-0020). M. Silva-Pereira acknowledges financial support from Fundação para a Ciência e a Tecnologia (2021.04898.BD). João Quelhas was affiliated with Banco de Portugal during most of the project. The views expressed herein are those of the authors and should not be attributed to any of the affiliated institutions. A working paper version was previously circulated with the title "A Temporary VAT Cut in Three Acts: Announcement, Implementation, and Reversal".

†Email: tiago.bernardino@iies.su.se
‡Email: ricardo.f.duquegabriel@frb.gov
§Email: joao.quelhas@su.se
¶Email: marcia.pereira@novasbe.pt

1 Introduction

Between 2020 and 2023, at least 21 countries enacted temporary value-added tax (VAT) reductions on food products to mitigate the adverse effects of rising living costs (Asquith 2024). However, the effectiveness of these measures has sparked controversy among economists and policymakers.

The canonical tax incidence model states that the economic burden of consumption taxes depends on the relative elasticities of supply and demand (Harberger 1962). This implies that the pass-through of a VAT change can range from 0% to more than 100% (Benzarti 2024). In most theoretical pricing models (e.g., constant markups), consumers are expected to bear the full burden of the tax since it does not change optimal pricing by firms, implying that VAT changes should be fully passed through to consumer prices. However, recent empirical literature on consumption tax incidence has found that VAT cuts are often not fully transmitted to consumer prices (Kosonen 2015; Fuest et al. 2024), their effects dissipate quickly (Crossley et al. 2014), and prices often rise to higher levels than before the reduction upon reversal (Benzarti et al. 2020). This seems to be regardless of whether the cut is temporary (Amores et al. 2023; Fuest et al. 2024; Forteza et al. 2024) or permanent (Benzarti et al. 2024), although pass-through estimates are higher in the first case.

In this paper, we estimate the pass-through of a temporary VAT cut to consumer prices, using a novel data set of daily online retail prices in Portugal. The policy affected 46 food products that had a tax cut from 6% to 0% between April 18, 2023, and January 4, 2024. We find a pass-through of 100% on impact, that persists for most of the policy duration and is symmetric upon reversal, suggesting that a temporary VAT cut can achieve a full, persistent, and symmetric pass-through to consumer prices in the context of supermarkets. These results appear to be consistent with the standard tax incidence models about VAT pass-through. However, we argue that the policy context, in particular the salience of the policy implementation, has played a major role in achieving such results.

We exploit a novel data set of daily online prices covering all food products sold in four major Portuguese supermarkets. It encompasses 43,283 items, of which 10,180 were affected by the VAT cut. Our empirical analysis is motivated by descriptive statistics comparing the price dynamics of the food items affected by the VAT cut (treated group) with the other food items that kept their VAT rate (control group). From our descriptive analysis, we draw three key findings. First, the prices of both treated and control groups followed similar trends before

the policy announcement. Second, we observe a striking break in the price series of the treated group on the day of implementation, resulting in a price gap that remains broadly constant until the official announcement of the reversal, after which the price gap closes completely. Third, we find that both treated and control groups exhibit similar patterns in the frequency and direction of price adjustments, except during the weeks surrounding the implementation and reversal of the policy.

To estimate the causal effects of the policy, we use a linear panel model with dynamic policy effects in an event-study setting. This approach is suitable for assessing the effect of such a policy change, given its clear announcement and implementation dates.

We document a sharp and substantial decrease in the relative price between treated and control items. On impact, prices were 5.66% lower, implying a pass-through estimate of 100%. This effect was persistent. When we extend the time window to assess the duration of the policy's impact, we find that the estimated pass-through remains close to 100% for several months, up until the announcement of the policy's termination. This indicates that the VAT cut was fully transmitted to consumer prices and that this transmission was long-lasting. The estimated treatment effect only begins to decline slightly during the final weeks of the policy.

We also study the reversal of the zero VAT policy separately, considering its official announcement as the starting point. We find that the price gap between the treatment and control groups begins to decrease slightly after the end of the policy announcement. Further, we uncover a symmetric pass-through to the implementation when the policy is reverted. When the policy ended, prices of the goods included in the VAT cut basket jumped by 5.98%, demonstrating symmetry of the pass-through between the start and end of the policy.

Leveraging the richness of the data, we then test the heterogeneity of the pass-through estimates. We run separate event-study regressions on different subsamples by product categories, brand (trademark versus white label), origin (domestic versus imported), supermarket size (big versus small), and prices before the policy change (below and above the median). We find that, when the policy was implemented and when the policy was reverted, the large majority of the categories had a treatment effect contained in the 95% confidence band of the average treatment effect. We interpret this as full pass-through of the VAT cut and its reversal across the different product dimensions.

Next, we assess whether the pass-through estimates are robust to other specifications using a comprehensive battery of tests. First, using aggregate data from the Harmonized Index

of Consumer Prices (HICP) for Portugal and Spain, we confirm that the estimated policy effect is also observed in monthly data used by the respective statistical offices to compute official inflation statistics. This finding corroborates that our results are not specific to online platforms and also apply to physical outlets, further validating the identification strategy. Second, we employ the synthetic difference-in-differences (SDiD) approach proposed by Arkhangelsky et al. (2021), which estimates time and item weights to make pre-trends similar between treatment and control groups. Third, given the critical importance of selecting an accurate control group, we explore alternative combinations of products: all products sold in supermarkets, only food and drinks, only nonfood products, and only those products with a 6% VAT rate. Finally, we test the robustness of the main results with other data treatments and outcome variables, such as prices without discounts or prices per unit. All estimates point toward the full and symmetric pass-through of the temporary VAT cut in Portugal for the 46 supermarket food products included in the VAT cut basket.

To better understand the potential mechanisms underlying our results, we discuss two potential explanatory factors: consumers' increased awareness of the policy amid high inflation and the decline of producer prices when the VAT cut occurred. The VAT cut was highly salient to consumers: It received extensive media coverage for a prolonged period, garnered significant public attention, and was tracked by the Portuguese Association for Consumer Protection and the main newspapers. The products with a VAT cut were identified through advertisements, banners, and stickers, and the high-inflation environment made consumers more attentive to price changes. Additionally, the policy's introduction coincided with a deflationary trend in producer-level food prices. This trend in supermarkets' input costs – not only for the treated goods but for control ones as well – may have facilitated the full pass-through observed at the retail level. Our results suggest that effective policy communication and the timing of the intervention relative to input price dynamics significantly influenced the policy's success.

Finally, we estimate the impact of the policy on the monthly inflation rate as measured by the HICP. Using the estimated impact that the policy had on treated food items in combination with the weight of this set of products in the official consumption basket, we estimate that the headline inflation rate fell by 0.68 percentage points (pp). This quantifiable impact on inflation highlights the potential of targeted fiscal policies as a tool for managing inflationary pressures.

Related Literature and Contribution. Our paper contributes to the literature on the incidence of VAT changes. We find a full pass-through of the tax cut in Portugal, with our estimates exceeding the upper bound of previous studies. In the context of food products, the range of pass-through estimates varies from 50% (Benzarti et al. 2024) to 100% (Gaarder 2019). Moreover, Buettner and Madzharova (2021) analyze VAT changes in the durable goods sector and found close to full but delayed pass-through into prices, suggesting that the nature of the product – whether essential or durable – can significantly affect the timing and extent of VAT pass-through. Other studies have explored the same phenomenon in industries such as restaurants (Harju and Kosonen 2014; Harju et al. 2018; Benzarti and Carloni 2019), hairdressing (Kosonen 2015), cinema (Arce and Antonio 2020), gasoline (Gautier et al. 2023; Montag et al. 2023), and online retailing (Fedoseeva and Van Droogenbroeck 2024), persistently finding pass-through rates below our estimates.

Our result is consistent with the canonical model of tax incidence, and resonates with older studies by Poterba (1996) and Besley and Rosen (1999), which found that sales tax hikes are fully passed on to customers, and, more recently, Gaarder (2019), who found that a permanent VAT cut in food items in Norway was fully shifted to consumer prices. Despite this, there is vast and growing evidence of departures from full pass-through and from symmetric incidence. Blundell (2009) notes that there are alternative theoretical models of product market competition that predict varying degrees of tax pass-through. Information asymmetry and the degree of competition, as demonstrated in Hindriks and Serse (2019) and Bellon et al. (2024), are some of the main explanations for the differing estimates. In our analysis, we argue that the salience that the VAT cut had for consumers in a high-inflation environment and producer price dynamics may be alternative explanations for the complete and persistent pass-through estimated. Alternatively, the asymmetric incidence observed in other studies may stem from firm or industry-specific factors that are not present in the specific context of supermarkets.

We also show that the pass-through of the temporary VAT cut in Portugal was symmetric. Contrary to the conventional theoretical predictions on symmetric incidence, Benzarti et al. (2020) demonstrate that prices respond more strongly to VAT increases than decreases using two exogenous changes in the tax rate for Finnish hairdressers and VAT changes in the European Union. Similarly, Politi and Mattos (2011) found a qualitatively similar result for tax changes in the Brazilian food market. Karadi and Reiff (2019) find that the frequency of

price changes is asymmetric: a greater number of goods experience price adjustments following a VAT increase compared to a VAT decrease. Our study complements these results by presenting evidence on a particular setting where symmetric incidence holds, consistent with standard models.[1]

The closest study to our analysis is Gaarder (2019), who examines a permanent VAT reduction on food in Norway and also documents a full pass-through. Our study differs in three key aspects: first, we analyze a temporary VAT cut and its reversal, allowing us to study symmetry in price adjustments; second, we use high-frequency daily data, rather than monthly price indices, providing a more granular view of price dynamics; and finally, we identify mechanisms behind the full pass-through, particularly consumer salience and producer price trends, which help explain why firms adjusted prices so consistently despite the tax cut being temporary.

Outline. The rest of the paper is organized as follows. In Section 2, we describe the background of the policy and detail its timeline. In Section 3, we present the data set of daily supermarket prices. In Section 4, we explain the empirical strategy adopted. Section 5 shows that the pass-through of the temporary VAT change in Portugal was complete, persistent, and symmetric. In Section 6, we discuss the mechanisms driving our results. Section 7 concludes.

2 Institutional Background

In November 2021, headline inflation in Portugal, as measured by the year-on-year percentage change of the HICP, exceeded the European Central Bank's reference target of 2% and embarked on an upward trend, peaking at 10.6% in October 2022. As in other European countries, food prices in Portugal increased more rapidly, in particular after the invasion of Ukraine (see Figure A.1 in the Appendix).

With food prices rising, public pressure mounted for the government to take action. This pressure intensified further when the Spanish government introduced a reduction in the VAT on a basket of essential goods at the end of 2022. Nevertheless, 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, as a consequence, on households' purchasing power because retailers would absorb a significant portion of the VAT reduction

[1] Related literature looks at the effects of tax changes on consumption and finds strong temporary effects but limited or no intertemporal substitution effects (e.g. Cashin and Unayama 2016; Baker et al. 2021).

in profit margins. On March 14, 2023, the Minister of Finance reinforced that no VAT change would lead to lower food inflation, as shown by the clippings in Figure A.2 in the Appendix.

On March 24, 2023, against the previously set expectations, the Portuguese government changed its stance, and the Minister of Finance announced a cut in the VAT to 0% on a selected list of 46 essential food products. The purpose of this policy was to fight the effects of inflation on households' purchasing power. The choice of the food products subject to the VAT cut was based on recommendations by the Directorate-General for Health, following the principles of a healthy and balanced Mediterranean diet and the most commonly consumed foods by the Portuguese population. It included, for example, bread, potatoes, pasta, and rice, and it had full coverage, meaning that it applied to all regions and was not targeted to any specific product dimension, such as low-priced items or white-label brands. Despite the selection of goods following a set of criteria, some appeared somewhat arbitrary. For instance, red beans were included, while white beans and black beans—which are substitutes—were not. The policy was enacted on April 18, 2023, for those 46 essential food items. All of them were previously taxed at a 6% rate, with the exception of vegetable oils, which were taxed at a 23% rate. The complete list of products in the VAT cut basket can be found in Table B.1 in the Appendix.

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 posted consumer prices, making it less salient to consumers during purchases. Firms collect the VAT from consumers and remit it to the tax authority, 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, this tax represented 21.3% of total government revenues in 2022 (9.4% of GDP), making it the most important source of tax revenue. There are three VAT regimes: (1) the standard rate of 23% applied to the majority of goods, (2) an intermediate rate of 13%, and (3) the reduced rate of 6% for certain essential goods.

Importantly, the policy attracted significant attention from the media and the general public. Consumer associations, journalists, and even some government officials closely monitored its implementation to assess whether supermarkets were raising posted prices above the original pretax prices.[2] Additionally, the labels of products included in the VAT cut indicated that

[2]The Portuguese Association for Consumer Protection, for example, built a price tracker of some products included in the VAT cut basket and regularly shared it in the media.

they were subject to a 0% rate, making this information highly salient to consumers (Figure A.3 in the Appendix).

The policy was originally announced as a temporary price relief measure that would last until the end of October 2023. However, in September, the government announced an extension until the end of the year, and the official ending date was only announced on October 27. On January 5, 2024, the policy was reverted, and the VAT rate applied to the targeted products returned to 6%. In summary, the timeline of the temporary VAT cut policy was as follows:

- March 24, 2023: Announcement of a temporary VAT cut for "essential products"
- March 27, 2023: Announcement of the list of products included in the VAT cut basket
- April 18, 2023: Implementation of the policy
- October 27, 2023: Announcement of the official ending date of the policy
- January 5, 2024: Reversal of the VAT cut on all the products

Given this timeline, we use the two announcements, one for the implementation and the other one for the reversal, as exogenous variations to study the dynamics of the pass-through of a VAT cut. These variations allow us to examine the dynamics of consumer prices throughout the entire policy life cycle. In the next section, we describe the data used for this analysis.

3 Data

3.1 Data Description

Our analysis uses a data set with supermarket daily prices (SDP) curated by Banco de Portugal Microdata Research Laboratory (BPLIM) (2024). This data set includes daily prices from the online stores of the main Portuguese retailers. Collectively, these account for more than half of the retail market share in the country. The information is collected using automated web-scraping algorithms on a daily basis.[3] For all items sold on each retailer's website, BPLIM stores information about their name, brand, units, capacity, bar code, and price. For the latter, both the posted price and the regular price (without discount) are collected, when available. Afterward, the posted price per unit is obtained using the number of units sold as a bundle.

The European Classification of Individual Consumption according to Purpose (ECOICOP) is used to classify each item into its category using the five-digit level of disaggregation. The

[3] In the last decade, the use of online prices has been increasing not only for measurement purposes but also for empirical research (Cavallo 2013; Cavallo and Rigobon 2016).

data set also includes a list of the products covered by the VAT cut using similar methods and guidance provided directly by the supermarkets, meaning that we are not required to perform this identification.

We focus on the period from January 2023 until the end of February 2024. During this time window, the data set encompasses around 63,000 items, defined as a product $\times$ supermarket combination. Several data-cleaning procedures are performed on this data set. Products that raised doubts about whether they were covered by the measure are excluded. Vegetable oils that had a decrease in the VAT rate from 23% to 0% are analyzed separately in Appendix A.3. Additionally, we exclude drinks, as these food items are not part of the treatment nor the control in the baseline specification presented above. Our final sample consists of 43,283 items, among which 10,180 (24%) were treated by having a reduction in the VAT rate from April 18, 2023, to January 4, 2024. Table B.2 in the Appendix details all of the steps.

We have an unbalanced panel, as some items are not always available or have missing price information.[4] Additionally, the items can be classified into different categories. Each item has an ECOICOP four-digit classification, which indicates the product category to which it belongs (e.g., bread and cereals, or fruit). We also classify items into trademark or white-label brands based on their brand. For the white-label classification, we identify items whose brand includes the name of the supermarket. Furthermore, we distinguish between imported and domestic items using the first three digits of the EAN bar code.[5] We classify larger supermarkets as those with a higher market share and smaller supermarkets as those with a smaller market share. Finally, we distinguish between high- and low-priced items using the median price of the treated items before the policy announcement. The number of treated and control items included in each of these categories is shown in Appendix B.2.

For our estimations, we use item weights which allows us to match our estimates following the structure of the HICP consumption basket and approximate the effect on inflation. The details on weighting are detailed in Appendix B.3.

[4] In Appendix C.4, we validate that our results hold even if we impose a balanced panel for the estimation.

[5] The EAN bar code corresponds to the European Article Number, in which the first three digits identify the country where the bar code was issued. We use this information to identify items produced domestically (EAN code starting with the digits 560) and those imported.

3.2 Descriptive Statistics

Figure 1 presents the price dynamics of the food items included in the VAT cut basket vis-à-vis the remaining food items. These groups will be considered as treatment and control, respectively, hereafter. The two plotted lines show the average of the daily price index between January 1, 2023, and January 31, 2024, normalized to 100 on the day before the policy was unexpectedly announced on March 23, 2023. The vertical dashed lines point toward the relevant moments under analysis: the implementation on April 18, 2023; the reversal on January 5, 2024; and the respective announcements of each VAT change. The items included in the treatment basket were taxed at the rate of 6% before the policy implementation and returned to this rate on the reversal day. The additional horizontal line shows the price level in case the VAT cut is fully transmitted. Four empirical patterns emerge in this figure.

Figure 1: Food Prices during the Temporary VAT Cut in Portugal

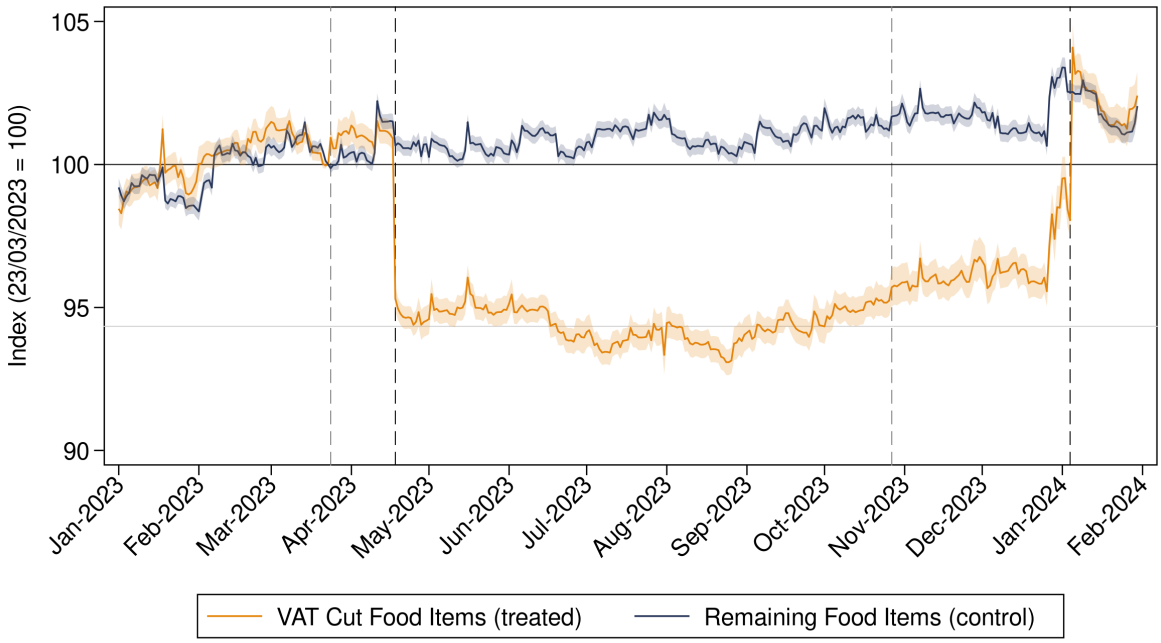

Notes: This figure shows the price dynamics of the food items included in the VAT cut basket (treated, in orange) and the remaining food items, excluding drinks (control, in blue). The data represent an average of the daily price index for each item from January 1, 2023, to January 31, 2024, normalized to 100 on the day before the unexpected announcement of the policy. The figure includes a band defined by two times the standard deviation around the average. The average and the standard deviation are computed using item-level price data for each day. The first dashed line indicates the announcement of the measure (March 24, 2023), the second dashed line marks the implementation day (April 18, 2023), the third dashed line indicates the official announcement of the reversal (October 27, 2023), and the fourth dashed line marks the reversal date (January 5, 2024). The horizontal solid line represents the price level corresponding to the full pass-through of the VAT cut policy. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

First, the food products that were not subject to the VAT cut seem to be a plausible control

group for the treated basket, as prices followed a similar trend before the announcement of the policy. This observation suggests that both groups of products were on parallel trends in the pre-reform period. Second, there was a striking break in the price series for the treated group when the policy was implemented. Third, the difference between the two groups of food products remained relatively constant until the announcement of the end of the policy at the end of October. After this announcement, the price gap between the two groups of products shrinks. Fourth, on the reversal day, the average price index of the treated products returned to the same level as that of the control group, showing a symmetric response of prices to this VAT change.

These observations indicate a full, persistent, and symmetric transmission of the VAT cut policy into consumer prices and motivate the empirical analysis that follows, in which we formally validate these findings. However, these observations also show that this type of data is highly volatile and subject to frequent changes. Therefore, we conduct the analysis using periods of five days.[6]

As a case study, we look at the price index of treated beans compared to non-treated beans (Figure A.4 in the Appendix). The price dynamics behave accordingly to what one would expect. First, both groups have a similar path before the implementation of the policy, and then a gap opens up between the two series, which is maintained until the day of the reversal when it returns to the previous price level.

One worry with our analysis may be that the treatment and control groups have different price-updating frequencies. To assuage that concern, Figure 2 illustrates the weekly frequency of price adjustments for the same baskets of items used above. An item is considered to have experienced a positive price change in a given week if its last price is higher than the last price of the previous week, and the opposite for a negative price change.[7] The left-hand-side panel displays the frequencies for goods included in the VAT cut basket, while the right-hand-side panel presents the frequencies for the other food items. Each graph differentiates between the direction of the price adjustments, with price increases depicted by orange bars and decreases by blue bars. This figure displays two main results.

First, both treated and control groups exhibit similar patterns in the frequency and direction of price adjustments, except for the weeks close to when the VAT was implemented and, subsequently, reverted. On average, approximately 20% of the goods experienced price

[6] In Appendix A.3, we conduct our analysis without any data aggregation. Results are unchanged.
[7] We obtain similar results using the item-level average price for each week.

Figure 2: Frequency of Positive and Negative Price Changes

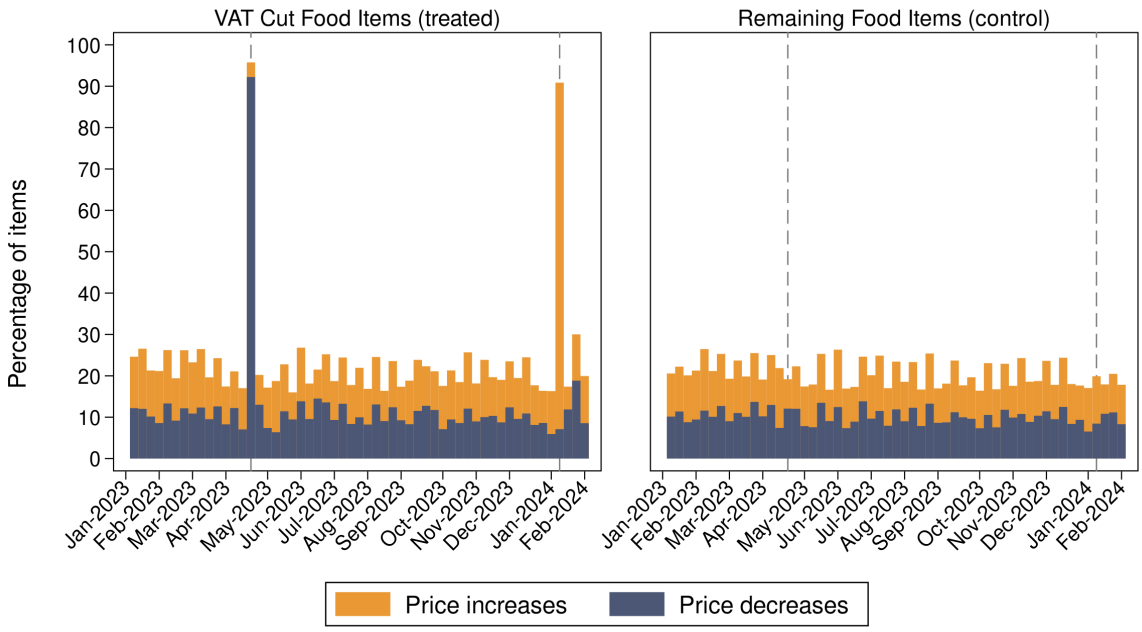

Notes: This figure shows the frequency of positive and negative price adjustments by week for the VAT cut food items (treated) and remaining food items (control). The stacked bars represent the percentage of items in a given week that experienced a change in the price level, with price increases shown in orange and price decreases in blue. An item is considered to have experienced a positive (negative) price change in a given week if its last price is higher (lower) than the last price of the previous week. The left-hand-side panel displays the frequencies for goods included in the VAT cut basket, while the right-hand-side panel presents the frequencies for the remaining food items. The first dashed line indicates the implementation day (week starting on April 18, 2023), and the second dashed line indicates the reversal date (week starting on January 5, 2024). *Source: Authors' calculations based on the BPLIM-SDP dataset.*

changes within the same week, with half of those having price increases and the remaining price decreases. This result reinforces the adequacy of our control group in what concerns the frequency of price adjustments. It also underscores the relevance of a thorough analysis centered around these two significant events, as the impact was concentrated around them.

To further validate the consistency of our data, we compute the monthly frequencies of price adjustments. In this case, we obtain an average monthly frequency of price changes of 30.5%. We exclude April 2023 and January 2024, as these months have abnormal pricing behavior. If we benchmark this sample against stylized features of consumer price-setting behavior in Portugal (Costa Dias et al. 2008; Martins and Quelhas 2024) or in other countries, such as the ones from the euro area (Gautier et al. 2024) or the U.S. (Nakamura and Steinsson 2008), we find similar figures for food.

Second, the total frequency of price changes is symmetric when the VAT decreases and when it increases back to the initial level. In the week when the VAT cut was implemented,

92% of the goods in the VAT cut basket experienced a price decrease. Conversely, during the week when the VAT was reinstated, about 83% of these goods saw a price increase. This small difference contrasts evidence put forward by Karadi and Reiff (2019) and Benzarti et al. (2020), who find that the frequency of price changes is asymmetric, particularly that a greater number of goods experience price adjustments following a VAT increase compared to a VAT decrease. For example, in Karadi and Reiff (2019), after the VAT increases, more than half of the products analyzed have a price increase, but after the VAT decrease, only 27% of the products have a price change. Similar differences show up in Benzarti et al. (2020).

Finally, we examine the distribution of price changes during the week following the VAT cut implementation, compared to the week before its announcement. Figure 3 presents this distribution: the left-hand panel shows the results for items included in the VAT cut basket, while the right-hand panel displays the corresponding distribution for other food items. Most items in the VAT cut basket experienced price reductions of 5% to 7%, indicating full pass-through. In contrast, most other food items did not show significant price changes.

Figure 3: Magnitude of Price Changes around the Temporary VAT Cut in Portugal

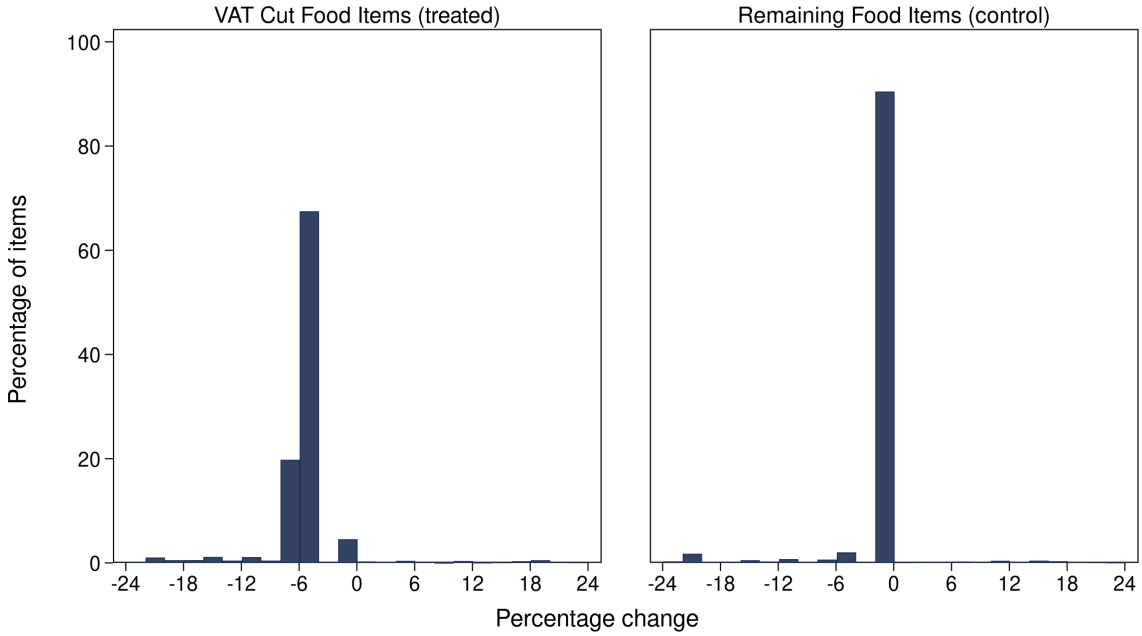

Notes: This figure shows the magnitude of price adjustments in the week following the implementation of the VAT cut compared with the week before its announcement for the VAT cut food items (treated) and the remaining food items (control). The bars represent the percentage of items in a given week that experienced a percentage change in the price level of the size of the respective bin, which corresponds to percentage changes. The left-hand-side panel displays the magnitudes for items included in the VAT cut basket, while the right-hand-side panel presents the magnitudes for the remaining food items. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

This pattern aligns with uniform pricing, where firms set prices as the product of marginal cost and a fixed markup (DellaVigna and Gentzkow 2019; Miravete et al. 2020). One possible explanation is that maintaining a constant markup simplifies administration, allowing firms to pass all cost changes – including variations in consumption taxes – directly onto consumer prices. Consequently, under this pricing rule, the incidence of VAT falls entirely on consumers, meaning that they bear the full burden of any VAT rate changes.

4 Empirical Strategy

We employ an event-study approach to formally estimate the effect of the VAT cut on the prices of affected food items, comparing these to the other food items that were not affected by the policy. Using the unique policy context described in Section 2, we analyze the consumer price dynamics before and after the announcement of the VAT reduction, as a starting point.

Our analysis uses the following linear panel model with dynamic policy effects:

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

where $P_{i,t}$ is the price index for each item i (a unique product $\times$ supermarket combination), observed in time t. α_i and γ_t denote item-specific and time-specific fixed effects, respectively. The variable $Z_{i,t-m}$ is the event indicator: It takes the value of 1 if the policy was announced at time $t - m$ and 0 otherwise. To facilitate the comparison relative to the policy's impact, the price index is normalized to 100 in the period immediately preceding the announcement, which enables a clear assessment of relative price movements post-announcement.[8]

A period t in Equation (1) corresponds to a five-day window, meaning that $P_{i,t}$ is the average price observed over those days. This methodological choice smooths out daily price volatility and aligns the data set to better capture the immediate and delayed effects of the policy. The five-day window is structured so that $m = 0$ coincides with the announcement and the subsequent windows capture policy effects, with $m = 5$ marking the onset of the implementation.[9]

[8] This approach follows the standard normalization practice in event studies, in which changes are measured relative to a defined pre-event period, ensuring that the effects observed are directly attributable to the policy. See, for example, Freyaldenhoven et al. (2021). The estimation uses the implementation by Freyaldenhoven et al. (2024).

[9] In Appendix B.4, we detail the dates matching to these periods and identify the periods in which each event occurred. Additional event studies where a period t corresponds to a single day are reported in Appendix A.3.

The parameters $\{\beta_m\}_{m=-G}^{M}$ quantify the dynamic effects of the policy, interpreted as the cumulative average treatment effects on the treated (ATT) at different time horizons m. Here, M and G define the range of periods pre- and post-event considered in the analysis, respectively. A positive β_m signals a price increase m periods after the policy announcement, whereas a negative value indicates a price decrease. While the coefficients for $m < 0$ estimate the anticipatory reactions to the policy announcement, $m > 0$ captures responses after the announcement.

The main identification hypothesis of our exercise is the parallel trends assumption, which posits that, in the absence of the VAT reduction, the price trajectories of the treated and control products would have been identical. Although directly testing this assumption is not feasible, we can examine the presence of parallel trends before the policy announcement and adjust our estimates accordingly to reflect any preexisting trends.

Our identification strategy also hinges on satisfying two critical conditions to ascertain the causal effects of the VAT changes. First, the Stable Unit Treatment Value Assumption (SUTVA) stipulates that there must be no spillover effects between the treated and control groups. Concerns regarding the potential inclusion of certain food products in the VAT cut basket (and the reverse) led us to include only those items for which the classification was unambiguous since the policy started, thereby upholding the integrity of SUTVA.

Additionally, a reduction in the VAT rate is expected to increase demand for treated goods, potentially causing the tax-inclusive prices of treated goods to increase relative to untreated goods. Such a shift could bias the point estimates downward, understating the true pass-through. This downward bias is likely to be particularly pronounced for close substitutes subject to differential treatment, rendering our estimates a lower bound of the actual pass-through.

Second, the assumption of no anticipation asserts that the policy was unexpected, meaning that neither consumers nor supermarkets could have adjusted their behaviors before the actual announcement of the policy changes. As discussed in Section 2, the Portuguese Finance Minister publicly rejected this policy several times, including only 10 days before its announcement. It is thus very unlikely that consumers and supermarkets reacted in advance.

Moreover, our model incorporates two sets of fixed effects to robustly estimate the policy impact. On the one hand, item fixed effects are used to control for unobservable attributes that do not vary over time, such as brand value or product expiry date of validity. Time fixed

effects, on the other hand, help account for external shocks or time trends that might affect all products simultaneously; these include macroeconomic conditions, seasonal variations, or general inflationary pressures. These fixed effects are crucial for isolating the effect of the VAT cut from other confounding influences. Standard errors are clustered at the item level.[10] We also normalize β_{-1} to zero, allowing us to interpret the plotted coefficients as estimated effects relative to the policy's impact. Finally, we carefully selected a control group consisting solely of other food items, excluding drinks, sold by retailers.

5 Full, Persistent, and Symmetric Pass-Through

5.1 Full Pass-Through

Figure 4 reports our main estimates of Equation (1). It displays the cumulative price dynamics, summarized by $\{\beta_m\}_{m=-G}^{M}$, after the announcement of the policy. The estimates reported in the figure correspond to the deviation from the extrapolated linear trend before the policy change, as in Dobkin et al. (2018). In the Appendix (Figure A.5), we show the estimated coefficients without trend correction. The shaded region in the figure corresponds to the pointwise 95% confidence interval, using standard errors clustered at the item level.

The coefficients fluctuate around zero in the days leading up to the announcement, which indicates no significant price changes in anticipation. Immediately after the announcement, the relative price between treated and control goods increased slightly by 0.81%. This increase is, nevertheless, short-lived. In fact, before the policy started, the relative price of treated goods returned to the same level as in the periods before the policy announcement.

The VAT cut started 25 days after the policy announcement. One can see that, immediately after the policy's implementation, there was a sharp and substantial decrease in prices. On impact, prices were 5.66% lower, corresponding to a full pass-through.[11]

We formally test for the presence of pretrends and find no evidence of their existence. Nonetheless, we apply the pretrend correction method proposed by Dobkin et al. (2018), which involves estimating the ATT using Equation (1) with a linear trend. This adjustment ensures that our coefficients of interest capture the precise quantitative effect of the policy change on consumer prices. However, this correction introduces a trade-off, as it increases the uncer-

[10] We also experimented with clustering standard errors at the ECOICOP level, and all results remain statistically significant at the same confidence level.

[11] A decrease in the VAT rate from 6% to 0%, corresponds to a $-0.06/1.06 \times 100 = -5.66\%$ change.

Figure 4: Pass-Through Estimates of the 2023 Temporary VAT Cut in Portugal

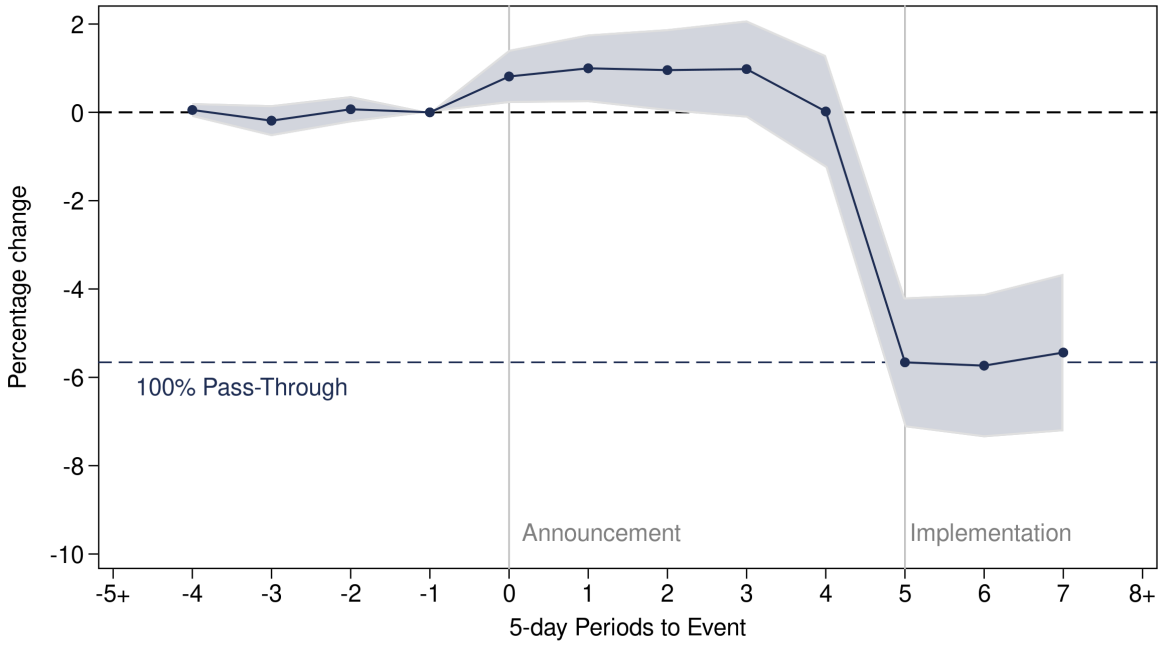

Notes: This figure shows the event-study estimates from Equation (1) after the policy announcement. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 7$. The estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). Item-level weights are used in the estimation. The shaded areas correspond to the 95% confidence interval with standard errors clustered at the item level. Similar figures without linear trend extrapolation and with daily data are in Appendix A.3. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

tainty of the point estimates. For comparison, Appendix A.3 presents the estimation results without the linear trend correction.

Overall, the event-study analysis provides compelling evidence that the temporary VAT cut policy had a tangible and immediate downward effect on consumer prices at the time of its introduction, as indicated by the narrow confidence intervals around the point estimates. These intervals are sufficiently precise to reject economically meaningful deviations from full pass-through, and even the lower bounds of our estimates remain consistent with the highest pass-through rates reported for food items in the literature.

While previous studies have examined the pass-through of tax changes, particularly tax cuts, almost none have found a full pass-through on impact (Benedek et al. 2020). One exception is the work by Gaarder (2019), who also finds complete pass-through of a permanent VAT cut on food in Norway. However, unlike the Norwegian setting, the VAT cut in Portugal was explicitly temporary. In Section 6, we discuss how the dynamics of producer prices and the salience of this policy for consumers may have affected the high pass-through observed.

In Appendix C, we show that the full pass-through result is robust to different assumptions, modeling, and data choices. One notable exercise adds external validity to our baseline by comparing Portuguese with Spanish price dynamics (see Appendix C.1). For this exercise, we use monthly data at the ECOICOP five-digit level, which allows us to also test if the results can be generalized beyond supermarket online stores. We also find a full pass-through in this case.

5.2 Persistent Pass-Through

To explore the persistency of the pass-through, we estimate the event-study model in Equation (1), allowing for a longer horizon, i.e. from the policy implementation until the end of the policy announcement. Figure 5 plots the estimated coefficients and the respective 95% confidence interval, using again standard errors clustered at the item level. In this analysis, contrarily to the previous estimates, we do not correct for the pre-trends, as in Dobkin et al. (2018). Consequently, the estimated coefficient on impact, which corresponds to the first point in the figure, is slightly different from the one reported in Figure 4.

Figure 5: Pass-Through Persistency of the 2023 Temporary VAT Cut in Portugal

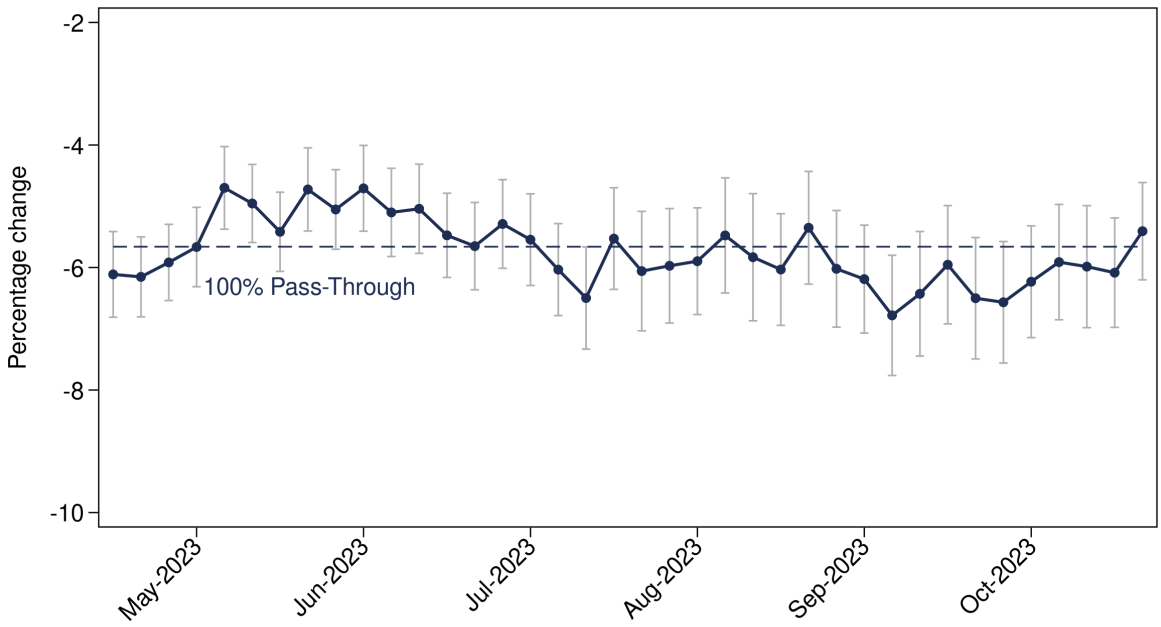

Notes: This figure shows the event-study estimates from Equation (1) between the implementation and the reversal of the policy. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 41$ without any pre-trends adjustment. Item-level weights are used in the estimation. The gray bars correspond to 95% confidence intervals with standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

The full pass-through estimated on impact remains high throughout the entire duration of the policy. Almost all the 95% confidence intervals of the estimated coefficients include the 100% pass-through line. This persistency in the full pass-through is an unusual result in the literature examining the effects of a temporary VAT cut on consumer prices.

Figure A.6 complements this analysis by showing how the magnitude of the price changes relative to the day before the announcement of the policy is distributed one, three, and six months after. It plots the distributions of the cumulative price changes of treated items, for each product i. These are fairly similar over time with a spike at the full pass-through level.

5.3 Symmetric Pass-Through

On October 27, 2023, the government announced the ending date of the policy to be on January 5, 2024. Even though there was an official announcement of the policy's end, it was known from the beginning that the policy was a temporary measure. Initially, it was supposed to last six months, but then the government extended it for three more months.

Figure 6 plots the event-study coefficients estimated using Equation (1) considering the announcement of the end of the policy as a reference date. The shaded region in the figure corresponds to the 95% confidence interval around the point estimates, using standard errors clustered at the item level. The five-day window is structured so that $m = 0$ coincides with the reversal announcement, and the subsequent windows capture post-announcement effects, with $m = 14$ marking the onset at the policy's reversal.

Before the reversal announcement, there were no significant price differences between the items included in the VAT cut basket and the remaining ones, which is consistent with the persistence of the full pass-through discussed above. After this announcement, the coefficients entered a slight upward trajectory. When the policy ended on January 5, 2024, the prices of the goods included in the VAT cut basket jumped by 5.98%, relative to the period before the reversal announcement.

This jump implies a full pass-through at the reversal, despite some announcement effect, closing the gap between the control and treated groups at the end of the policy. This indicates a symmetry of the pass-through between the start and end of the policy. We show the robustness of this result for a comprehensive battery of tests in Appendix C.

This symmetric effect is noteworthy. In a recent influential paper, Benzarti et al. (2020) show that consumption taxes have an asymmetric incidence, arguing that prices respond sig-

Figure 6: Pass-Through Estimates of the 2023 Temporary VAT Cut Reversal in Portugal

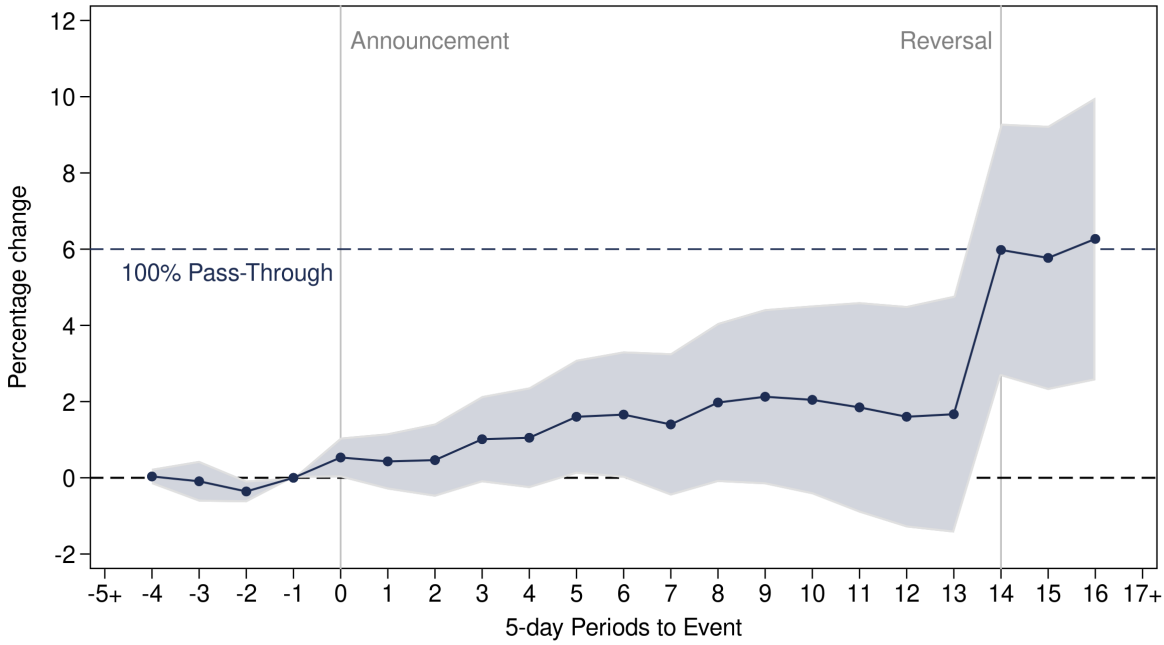

Notes: This figure shows the event-study estimates from Equation (1) after the official announcement of the policy end. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 16$. The estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator (Dobkin et al. 2018). Item-level weights are used in the estimation. The shaded areas correspond to the 95% confidence interval with standard errors clustered at the item level. Similar figures without linear trend extrapolation and with daily data are in Appendix A.3. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

nificantly more to increases than to decreases in VATs. Our findings provide evidence that the result of symmetry predicted by standard models can still hold. In Section 6, we discuss this result further, providing reasons for the observed symmetry.

5.4 Heterogeneity Analysis

Treatment effects may depend on product characteristics. We now extend our analysis to study differences in the pass-through across multiple dimensions. To do that, we create different subsamples according to various product characteristics. We examine whether any heterogeneity in treatment effect occurs across product types (using the four-digit ECOICOP classification), brands (trademark versus white label), origins (domestic versus imported items), supermarket size (big versus small), and prices before the policy change (below and above the median). Appendix B.2 reports the number of items in each category.

We estimate the coefficients in Equation (1), restricting the treatment group to products with specific characteristics, using all food items not included in the VAT cut basket as con-

trol, as in the baseline exercise. Figure 7 shows the estimated coefficients along the different dimensions. We look at the price dynamics at both the implementation, on the left side, and the reversal of the policy, on the right side.

Figure 7: Heterogeneous Treatment Effects Analysis

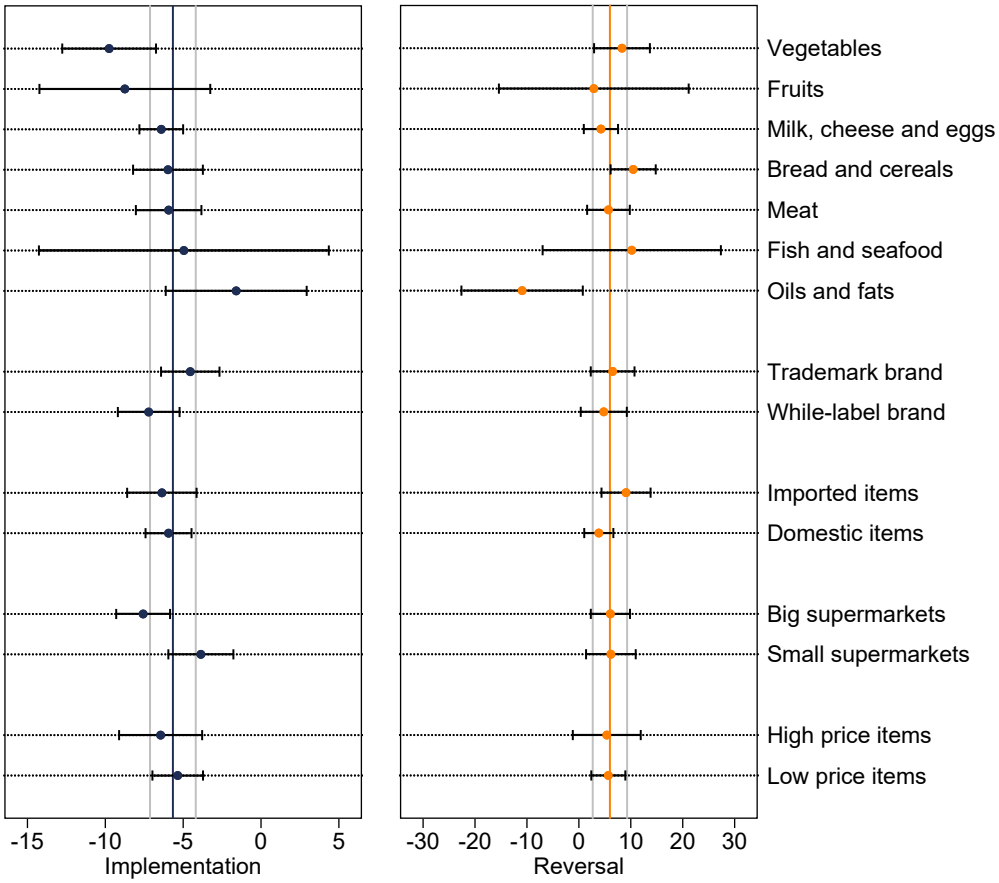

Notes: This figure shows the heterogeneous treatment effects along different product dimensions in the implementation (left-hand-side panel) and the reversal (right-hand-side panel) of the policy. The treatment effect corresponds to the event-study estimate from Equation (1) on the respective five-day period of the implementation (blue dots) and the reversal (orange dots), using subsamples created according to the different dimensions on the right: product categories (ECOICOP four-digit level), brand, origin, size of the supermarket, and price before the policy change. Product categories are ordered according to the size of the implementation coefficients. The vertical bars in each panel correspond to the aggregate coefficient for each of the two moments, reported in Figures 4 and 6, along with the respective 95% confidence interval lower and upper bounds in gray. The range of periods pre- and post-event considered in each analysis is the same as in the previous exercises. The estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator (Dobkin et al. 2018). The black solid lines correspond to the 95% confidence interval with standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

The first set of coefficients, at the top of the figure, shows the price change across different food categories. On the left, we see that, at the implementation stage of the policy, for the majority of food categories, the estimated coefficients are within the 95% confidence band of the average treatment effect (represented by the gray vertical lines). This indicates that there

was a full pass-through of the VAT cut across multiple product dimensions. The vegetables category is the exception, exhibiting a pass-through above 100%. A possible reason for this is the decline in producer prices around the policy implementation for the products in this category (Figure A.13). In Section 6, we discuss in detail how falling producer prices can lead to price changes beyond the full pass-through result.

On the right, we observe a similar pattern following the policy reversal. For most product categories, the estimated treatment effect is within the 95% confidence band of the average treatment effect (represented by the gray vertical lines). This indicates that the average full pass-through on reversal also happened across almost all food categories. A notorious exception is the oils and fats category. In Figure A.9 in the Appendix, we plot the event study without the linear trend correction. The figure shows that before the announcement of the policy reversal, the price of this category was on an upward trajectory in comparison with the control group.[12] This upward trajectory is justified by a poor harvest year for olives, which contributed to the surge in the product prices of this category, particularly olive oil.

Regarding the other dimensions reported in the bottom part of both subfigures, we find that the estimated effect is contained in the 95% confidence band of the average treatment effect. This suggests that the pass-through was not different across these other product characteristics. In Figure A.10 in the Appendix, we complement this heterogeneity analysis by presenting the total net treatment effect between the policy announcement and its reversal for the same categories considered above. Despite some caveats in this analysis, as discussed in the Appendix, the broader picture reveals a symmetric pass-through across the different categories, with most dimensions exhibiting a net effect of approximately zero.

Overall, this likely reflects the relatively short duration and temporary nature of the policy and Portugal's specific economic context. Factors such as market competition, consumer behavior, and awareness about the VAT changes could play a crucial role in shaping these outcomes, as argued in Bellon et al. (2024). We discuss these aspects in more detail in Section 6.

[12] Since we apply the Dobkin et al. (2018) trend correction, the estimated coefficient for this category becomes smaller than the average treatment effect.

5.5 Robustness Exercises

Next, we summarize the different robustness exercises we perform to validate these results. Further details of the data, methods, and results can be found in Appendix C.

Identification Strategy and External Validity. Our baseline identification strategy relies on a comparison between food items that retained their VAT rate and those that experienced a rate reduction, both sold on the online stores from which the data are web scraped. As an alternative identification strategy, we consider the price dynamics of food items across different countries. For this option, we collect aggregate HICP data for the treated products, using the ECOICOP five-digit classification, in Portugal and Spain. These data are sourced from Eurostat and collected by the national statistical offices.[13] This means that the data are standardized, allowing us to compare the evolution of each product category over time in the two countries. Moreover, the prices underlying these data are representative of both online and physical retail stores, providing external validity to our findings. In Appendix C.1, we estimate the pass-through of the VAT cut. Our results indicate that the change in the VAT was fully transmitted to consumer prices at the implementation stage and remained high throughout the policy duration.

Synthetic Difference-in-Differences. The difference-in-differences (DiD) identification strategy relies on the hypothesis that the treated and control groups exhibit parallel trends before the event. In our baseline results, we adjust for minor deviations in pre-trends to accurately interpret the coefficients. Alternatively, we employ the SDiD approach proposed by Arkhangelsky et al. (2021), which estimates time and item weights to ensure the parallel pre-trends between control and treatment groups. In Appendix C.2, we describe in detail this method and present the results obtained, showing that prices had a decrease between -6.97% and -5.85%, upon implementation and, subsequently, had an increase between 6.52% and 7.86%, upon reversal, using as reference the respective preannouncement days and a 95% confidence interval. The estimates are consistent with the baseline results.

Control Group. The selection of products included in the control group is critical for accurately estimating the ATT, as it must include goods that are comparable with those in the VAT cut basket. In our preferred specification, we focus exclusively on food products, excluding

[13] In Appendix B.5, we describe the data sources in more detail.

drinks. In Appendix C.3, we explore the pass-through dynamics of the VAT cut using four alternative control groups. First, we consider all products sold in the supermarket, both food and nonfood. Second, we restrict the control group to food and drink products. Third, we include only nonfood products to eliminate potential substitution effects among treated and control items. Fourth, we use products that have a reduced tax rate of 6%, matching those in the treatment group. Although the estimated coefficients vary, the 95% confidence interval includes the full pass-through point estimate at the implementation and at the reversal, corroborating the robustness of our finding of a full and symmetric pass-through.

Data-Cleaning Methods. The data set comprises 43,283 items over 14 months. The panel is unbalanced because not all products are available for sale every day of the year. These gaps may arise from product unavailability or data collection errors from the web-scraping algorithms. In our baseline specification, we use the unbalanced data set without any additional further cleaning or imputation. In Appendix C.4, we test the completeness and symmetry of the VAT cut pass-through using two methods to impute missing data. The first method carries forward the price from the last available day for a maximum of one week. The second method carries forward the price without any time limitation on the missing gap. We also examine a more stringent scenario, retaining only products with prices observed for all periods under analysis, resulting in a balanced panel of 8,933 items. When prices are imputed without time restriction, the estimated coefficient at implementation falls below the full pass-through. In the other scenarios, the confidence interval consistently includes the 100% pass-through.

Outcome Variable. Another set of robustness exercises we conduct relates to the choice of the dependent variable. In our baseline specification, the dependent variable is the posted price as labeled, which includes the VAT and any applicable sales or discounts. In Appendix C.5, we estimate Equation (1) using two alternative dependent variables — the regular price, i.e., the price before sales, and the posted price per unit — to account for potential shrinkflation. We show that the results consistently indicate a full and symmetric pass-through.

6 Discussion

Our results show that the temporary VAT cut on a subset of food items was fully passed through to consumer prices, with the effect remaining persistent and symmetric upon rever-

sal. These findings align closely with predictions from standard models of tax incidence but stand in contrast to a growing body of recent empirical evidence documenting incomplete and asymmetric VAT pass-through. In what follows, we explore how the specific macroeconomic and institutional features of this policy episode may account for these outcomes and help bridge the gap between canonical theoretical predictions and recent empirical findings.

One potential explanation for obtaining a result aligned with standard models but at odds with these recent studies is that there may be firm or industry-specific factors present in those divergent cases that are absent in the context of supermarkets. For example, Harju et al. (2018), who examine VAT reductions in the case of restaurants, find that independent restaurants did not pass through any change to consumer prices, while chain restaurants fully passed the tax cut to consumers. This dichotomy between independent versus chain restaurants illustrates one scenario where firm-level factors can justify our result.

On the other hand, there might also be factors that are present in our context that could motivate the complete and symmetric pass-through, that do not hold for those anomalous cases. As previously mentioned, administrative costs associated with managing a large number of products provide a rationale for pricing at a constant markup, which would be consistent with a complete and symmetric pass-through and is specific to this industry. In industries where there are fewer products or services, this reasoning may not hold, thus justifying different pricing strategies that translate into incomplete and asymmetric pass-through.

Besides this, there are other features of this policy's setting that could be important drivers of the obtained result. In the rest of the section, we discuss more in-depth two other mechanisms, producer price dynamics and policy salience, for which we have suggestive evidence. We also discuss at the end of the section how the obtained results compare with the initially defined policy objectives.

6.1 Mechanisms

Deflation in Producer Prices. The time of the announcement and implementation of the policy coincided with producer prices embarking on a deflationary path. A fall in costs around the time of the policy means that retailers can fully pass through the VAT cut to consumer prices while keeping or even increasing their margins. It is important to note that deflation in producer prices is not exclusive to treated goods, being common to the vast majority of the goods for which we have producer price data. This fact means that retailers' margins on some

of the control items could be increasing, providing further incentive to pass the VAT cut to consumer prices on treated goods.

To test this idea, we conduct an event study on weekly product-level price data in agricultural wholesale markets in the spirit of the analysis for consumer prices, following the methodology described in Section 4.[14] We include both treated and control items because declining producer prices for goods not affected by the VAT cut additionally reduce retailers' average input costs. We also use a fixed-effects estimator on time and item, where the latter is a product-market combination. Our base period is the week preceding the announcement of the policy. Our window ranges from three weeks before the announcement until seven weeks after the announcement, with the policy being implemented three weeks after the announcement. Standard errors are clustered at the item level. The event study for agricultural wholesale prices departs from the empirical strategy for consumer prices in two ways: We do not use a control group due to insufficient available data on product categories, and we do not use trend correction.

Figure 8 presents the results of the event study for agricultural wholesale prices described above. Prices were on a downward trend even before the announcement, with the decline relative to the period preceding the announcement being statistically significant for the entire period after the event. This finding provides further evidence of the falling producer prices potentially motivating the full pass-through but also allowing for it to persist over time. Figure A.13 in the Appendix presents the producer price index for the total and four main categories of agricultural wholesale prices— meat, fruit, vegetables, and milk—which have the biggest weight in the consumer's basket. The total index showcases the same result, as we observe a decline in producer prices starting some weeks before the announcement of the policy until mid-May. In the case of meat, the fall begins later toward the end of April. The same pattern for the total is also observed for milk and, more starkly, for vegetables.

Comparing these results with the heterogeneous effects in Section 5.4, we can see that, at implementation, goods such as vegetables had a higher than full pass-through, coinciding with producer prices falling at the same time the policy was announced and implemented. As such, declining producer prices could be a mechanism for the unusually high pass-through we observe, allowing supermarkets to keep their margins or even increase them.

We complement this analysis by looking at the publicly available monthly industrial pro-

[14] We describe the data set in more detail in Section B.5.

Figure 8: Agricultural Wholesale Producer Prices

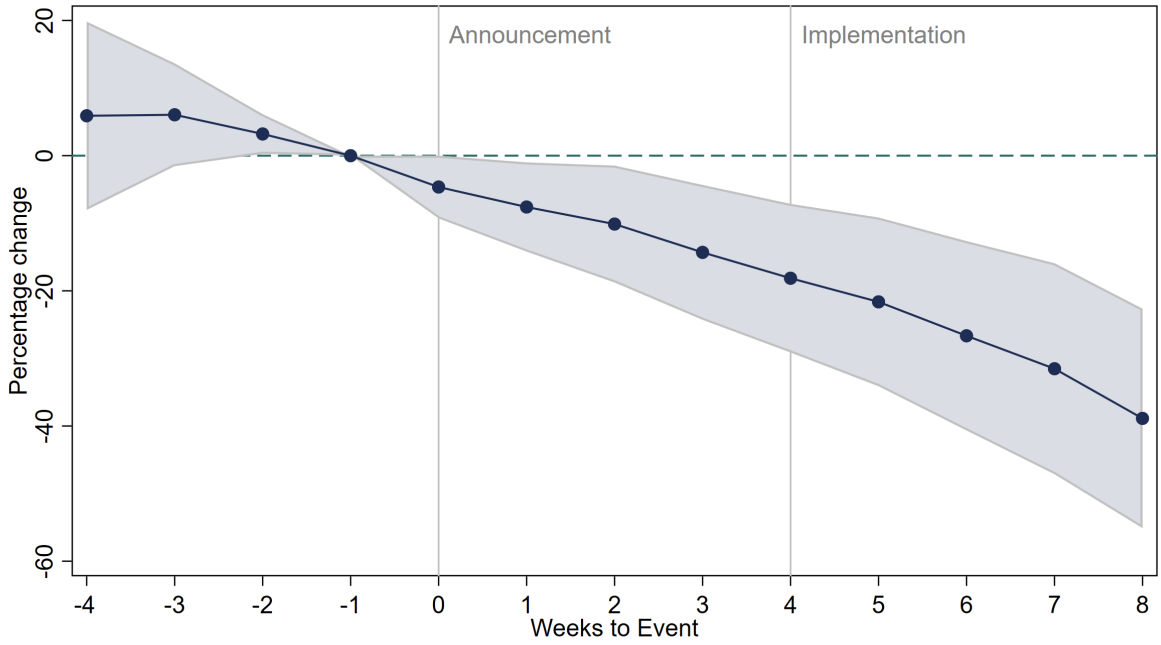

Notes: This figure shows the event-study estimates from Equation (1), with the exception that we do not consider a control group. We consider as a base period the week preceding the policy announcement, starting on March 20, 2023. The range of periods pre- and post-event considered in the analysis is $M = 3$ and $G = 7$. The shaded areas correspond to 95% confidence intervals with standard errors clustered at the product $\times$ market combination. *Source: Authors' calculations based on the price data of the Agricultural Markets Information System of the Portuguese Planning, Policy and General Administration Office.*

duction price index for manufacturing food industries by category, published by Statistics Portugal. This decrease in producer prices is also present when we look at this series, as presented in Figure A.14 of the Appendix. We can see that the total slightly declines from March 2023 to April 2023 and more afterward. Similar to wholesale agricultural prices, the fall is more pronounced for fruits and vegetables.

The deflation in producer prices may account for the full pass-through observed at the time of policy implementation. However, it falls short in explaining the symmetry in pass-through upon reversal, suggesting that additional mechanisms are at play. We argue that the institutional features of the policy, particularly the high degree of public awareness and salience, played a central role in shaping the response of economic agents. In what follows, we discuss how these features can account for both the completeness and the symmetry of the observed pass-through.

Salience. The VAT cut policy was extremely salient to consumers in several ways. First, as mentioned in Section 2, the discussion around the policy gathered significant media coverage

and public attention. The rapid escalation of food prices, anecdotal claims of price gouging by supermarkets, pressure from several stakeholders for a VAT cut, and denial from members of the government that such a policy would be enacted are some of the factors that contributed to the elevated public attention.

The degree of media attention and public scrutiny is another potential mechanism that incentivized food retailers to pass through the VAT cut to consumers. Throughout the policy duration, the Portuguese Association for Consumer Protection built a price tracker of some products included in the VAT zero basket, which was regularly shared in the media. Although this association does not have any legal power to, for example, impose fines, it contributed to increase consumers' awareness about the policy. The search intensity for the terms "IVA" and "Imposto sobre o valor acrescentado" (VAT and value-added tax, respectively, in Portuguese) on Google's search engine from January 2023 until February 2024 was very high, comparable with the search intensity for "Taylor Swift" in the same period, when the artist announced two concerts in Lisbon (see Figure A.11 in the Appendix). This provides suggestive evidence of the high level of public attention to the policy.

Second, upon implementation, supermarkets heavily publicized the VAT cut policy. Banners were saying "IVA 0" next to the shelves with products in the VAT cut basket, and several of the items in the basket had stickers also indicating the VAT exemption (see Figure A.3 in the Appendix). This publicity was a common practice across retailers, both in physical locations and online, making it much more salient to consumers which goods were included in the policy and which were excluded. The power of salience as a mechanism to explain the full and persistent pass-through of this temporary VAT cut is supported by Chetty et al. (2009), who provide evidence that consumers underreact to taxes that are not salient and that "the demand curve becomes more inelastic when individuals are inattentive". In our case, attention could then imply a higher tax elasticity of demand, thus providing further incentive for retailers to pass through the tax cut. Ramey (2021) also argues that changes in VAT and sales taxes are salient and understandable to the average consumer.

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

Moreover, we find evidence that consumers incorporated the VAT cut into their expecta-

tions. Following the policy announcement, the gap between the euro area's and Portugal's inflation expectations narrowed by nearly 2 pp. When the policy's official end was announced in late October 2023, this gap widened again (Figure A.12 in the Appendix). This development demonstrates that the policy's salience was strong enough to alter expectations. Our findings corroborate Bachmann et al. (2021), who show how a temporary VAT cut can be used to influence consumer expectations.

Other Mechanisms. Other mechanisms could play an important role in our result. Bajo-Buenestado and Borrella-Mas (2022) document that the pass-through of a tax on prices is about 38% higher in vertically integrated markets compared to white-label branded products. In Figure 7, we qualitatively validate this result, showing that white-label products have a higher pass-through, even though the difference is not statistically significant. A large share of the items in our data is white-brand labeled (19.5%), which could have contributed to the high pass-through we estimate. Fuest et al. (2024) also provide evidence of heterogeneity along market integration between trademark brand and white-label brand items.

A final set of reasons is related to the institutional setting of this particular policy. Before the announcement of the policy, 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 to ensure that these groups were committed to reducing and stabilizing prices. This type of agreement could create further incentives for agents across the supply chain to fully pass the VAT cut on to consumers. Furthermore, the government had instructed food safety inspectors from the Economic and Food Safety Authority to monitor prices, with particular emphasis on price speculation. During the nine months of the policy, the authority inspected approximately 2,000 economic operators and identified 174 cases of noncompliance with the VAT cut policy.

6.2 Policy Objectives

As mentioned in Section 2, the VAT cut policy had the explicit goal of mitigating the effects of inflation on household income and directly lowering inflation. Our results show that the pass-through of the policy to consumer prices was complete and persistent. Given the extent of this policy, we conclude that it was effective in lowering consumer prices. However, one could argue that there were better policy alternatives to fulfill this goal (e.g., direct transfers

to poorer households).[15]

According to the Portuguese Directorate-General for Budget, the policy implied a total loss of tax revenue of 521 million Euros, which corresponds to 2.27% of total value-added tax revenue and 0.89% of total tax revenue in 2023. Other policies were implemented during 2023 to counter the welfare costs of inflation, such as direct transfers. The total expenditure on direct transfers to households – targeted at vulnerable households, families with children, and retirees – amounted to 578 million euros, which is 11% more than the tax revenue lost due to the VAT cut policy.[16] This prompts us to consider the opportunity cost of the VAT cut and raises the question of the potential welfare implications of reallocating the lost tax revenue to alternative measures, such as direct transfers to households.

To evaluate the effectiveness of the policy in reducing inflation, we can approximate the policy's direct impact on aggregate inflation by combining the estimated price changes with each category's weight in the representative consumer's consumption basket. The weight of the goods included in the VAT cut basket was 12% in 2023.[17] This implies that the effect of the temporary VAT cut on the aggregate inflation rate is 0.68 pp. Note that this estimate represents the direct effect on inflation. It does not consider changes in the consumption patterns that can affect the weight in the HICP of VAT cut products nor other general equilibrium effects.

7 Conclusion

In this paper, we investigate the pass-through of a temporary cut in the VAT to consumer prices, using a novel data set of daily online retail prices in Portugal that covers the full universe of food products sold in supermarkets. The decrease in the VAT affected a subset of food items starting on April 18, 2023, and lasted for approximately nine months. We analyze the consumer price dynamics across the complete policy lifetime.

We find that the pass-through of the VAT cut to consumer prices was complete, persistent,

[15] The lack of access to granular data on consumption does not allow us to look into the heterogeneous effects that this policy could have had on households. Benzarti et al. (2024) analyze the distributional effects of a VAT cut and conclude that such a policy is more beneficial to poorer than richer households by a factor of 3. Similarly, Gaarder (2019) finds that a permanent VAT reduction on food in Norway led to full pass-through and had progressive distributional effects, benefiting lower-income households more than proportionally.

[16] The figures presented in this paragraph related to the government's fiscal revenue and spending are available in the summary of the government's budget execution, published every year-end by the Portuguese Directorate-General for Budget (*Direção-Geral do Orçamento*), titled *Síntese da Execução Orçamental de dezembro de 2023*.

[17] This number is computed based on the weight that each food category has on the HICP and the percentage of products within each category with a VAT cut. Both are provided by Statistics Portugal.

and symmetric. When the policy started, prices fell by 5.66%, which corresponds to a full pass-through. Despite recent burgeoning evidence on the incomplete pass-through and asymmetric incidence of consumption tax changes, our results are consistent with baseline knowledge on VAT incidence. We also find that this high pass-through is sustained until the government announces the end of the policy. Furthermore, we provide evidence that when the policy was reversed, the price of the goods included in the VAT cut basket returned to the trend of the other food items (the control group), exhibiting symmetry between the implementation and the reversal. These results are homogeneous across multiple dimensions and robust to several data and modeling choices.

Collectively, these findings not only shed light on the direct impacts of VAT adjustments on prices but also highlight the critical importance of timing, communication, market expectations, and institutional setting in maximizing the efficacy of public policies aimed at managing inflation. Investigating further 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

Amores, A., S. Barrios, R. Speitmann, and D. Stoehlker (2023). Price Effects of Temporary VAT Rate Cuts: Evidence from Spanish Supermarkets. *European Commission*.

Arce, I. and M. Antonio (2020). Una Evaluación sobre los Efectos de una Reducción del Tipo de Gravamen del IVA para los Bienes y Servicios Culturales. *Instituto de Estudios Fiscales*.

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

Asquith, R. (2024). Past Value-Added Tax and Goods and Services Tax changes. *VAT Calc https://www.vatcalc.com/global/global-2022-vat-gst-changes/*, (accessed: 16.07.2024).

Bachmann, R., B. Born, O. Goldfayn-Frank, G. Kocharkov, R. Luetticke, and M. Weber (2021). A Temporary VAT Cut as Unconventional Fiscal Policy.

Bajo-Buenestado, R. and M. Borrella-Mas (2022). The Heterogeneous Tax Pass-Through Under Different Vertical Relationships. *The Economic Journal 132*(645), 1684–1708.

Baker, S. R., S. Johnson, and L. Kueng (2021). Shopping for Lower Sales Tax Rates. *American Economic Journal: Macroeconomics 13*(3), 209–250.

Banco de Portugal Microdata Research Laboratory (BPLIM) (2024). Supermarket Daily Prices Database. Extraction: March 2025. Version: V1. BANCO DE PORTUGAL. Dataset.

Bellon, M., A. Copestake, and W. Zhang (2024). Supply and Demand Determinants of Heterogeneous VAT Pass-Through. *SSRN*.

Benedek, D., R. De Mooij, M. Keen, and P. Wingender (2020). Varieties of VAT Pass-Through. *International Tax and Public Finance 27*, 890–930.

Benzarti, Y. (2024). Tax Incidence Anomalies. *NBER, WP 32819*.

Benzarti, Y. and D. Carloni (2019). Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France. *American Economic Journal: Economic Policy 11*(1), 38–63.

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), 4438–74.

Benzarti, Y., S. Garriga, and D. Tortarolo (2024). Can VAT Cuts and Anti-Profiteering Measures Dampen the Effects of Food Price Inflation? *NBER, WP 32241*.

Besley, T. and H. 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.

Blundell, R. (2009). Assessing the Temporary VAT Cut Policy in the UK. *Fiscal Studies 30*(1).

Buettner, T. and B. Madzharova (2021). Unit Sales and Price Effects of Preannounced Consumption Tax Reforms: Micro-level Evidence from European VAT. *American Economic Journal: Economic Policy 13*(3), 103–134.

Cashin, D. and T. Unayama (2016). Measuring Intertemporal Substitution in Consumption: Evidence from a VAT Increase in Japan. *Review of Economics and Statistics 98*(2), 285–297.

Cavallo, A. (2013). Online and Official Price Indexes: Measuring Argentina's inflation. *Journal of Monetary Economics 60*(2), 152–165.

Cavallo, A. and R. Rigobon (2016). The Billion Prices Project: Using Online Prices for Measurement and Research. *Journal of Economic Perspectives 30*(2), 151–178.

Chetty, R., A. Looney, and K. Kroft (2009). Salience and Taxation: Theory and evidence. *American Economic Review 99*(4), 1145–1177.

Clarke, D., D. Pailañir, S. Athey, and G. Imbens (2023). Synthetic Difference-in-Differences Estimation. *IZA, Discussion Paper 15907*.

Costa Dias, M., D. Dias, and P. Duarte Neves (2008). Stylised Features of Consumer Price Setting Behaviour in Portugal: 1992–2001. *Portuguese Economic Journal 7*, 75–99.

Crossley, T. F., H. W. Low, and C. Sleeman (2014). Using a Temporary Indirect Tax Cut as a Fiscal Stimulus: Evidence from the UK. *IFS, WP W14/16*.

DellaVigna, S. and M. Gentzkow (2019). Uniform Pricing in U.S. Retail Chains. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Dobkin, C., A. Finkelstein, R. Kluender, and M. Notowidigdo (2018). The Economic Consequences of Hospital Admissions. *American Economic Review 108*(2), 308–352.

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 (2024). Analysing the VAT Cut Pass-Through in Spain using Web-scraped Supermarket Data and Machine Learning. *Banco de España, WP 17*.

Freyaldenhoven, S., C. Hansen, J. Pérez, and J. Shapiro (2021). Visualization, Identification, and Estimation in the Linear Panel Event-Study Design. *NBER, WP 29170*.

Freyaldenhoven, S., C. Hansen, J. Pérez, J. Shapiro, and C. C. (2024). xtevent: Estimation and Visualization in the Linear Panel Event-Study Design. *The Stata Journal*, forthcoming.

Fuest, C., F. Neumeier, and D. Stöhlker (2024). The Pass-Through of Temporary VAT Rate Cuts: Evidence from German Supermarket Retail. *International Tax and Public Finance*, 1–47.

Gaarder, I. (2019). Incidence and Distributional Effects of Value-Added Taxes. *The Economic Journal 129*(618), 853–876.

Gautier, E., C. Conflitti, R. Faber, B. Fabo, L. Fadejeva, V. Jouvanceau, J. Menz, T. Messner, P. Petroulas, P. Roldan-Blanco, F. Rumler, S. Santoro, E. Wieland, and H. Zimmer (2024). New Facts on Consumer Price Rigidity in the Euro Area. *American Economic Journal: Macroeconomics*, forthcoming.

Gautier, E., M. Marx, and P. Vertier (2023). How Do Gasoline Prices Respond To a Cost Shock? *Journal of Political Economy: Macroeconomics 1*(4), 707–741.

Harberger, A. C. (1962). The Incidence of the Corporation Income Tax. *Journal of Political Economy 70*(3), 215–240.

Harju, J. and T. Kosonen (2014). The Inefficiency of Reduced VAT Rates: Evidence from Restaurant Industry. *International Institute of Public Finance*.

Harju, J., T. Kosonen, and O. N. Skans (2018). Firm Types, Price-Setting Strategies, and Consumption-Tax Incidence. *Journal of Public Economics 165*, 48–72.

Hindriks, J. and V. Serse (2019). Heterogeneity in the Tax Pass-Through to Spirit Retail Prices: Evidence from Belgium. *Journal of Public Economics 176*, 142–160.

Karadi, P. and A. Reiff (2019). Menu Costs, Aggregate Fluctuations, and Large Shocks. *American Economic Journal: Macroeconomics 11*(3), 111–146.

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.

Martins, F. and J. Quelhas (2024). Consumer Price-Setting Behaviour: Evidence from Food CPI Microdata.

Miravete, E. J., K. Seim, and J. Thurk (2020). One Markup to Rule them All: Taxation by Liquor Pricing Regulation. *American Economic Journal: Microeconomics 12*(1), 1–41.

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, WP 337*.

Nakamura, E. and J. Steinsson (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Pfäuti, O. (2023). The Inflation Attention Threshold and Inflation Surges. *Working Paper*.

Politi, R. and E. Mattos (2011). Ad-valorem Tax Incidence and After-tax Price Adjustments: Evidence From Brazilian Basic Basket Food. *Canadian Journal of Economics 44*(4), 1438–70.

Poterba, J. (1996). Retail Price Reactions to Changes in State and Local Sales Taxes. *National Tax Journal 49*(2), 165–176.

Ramey, V. (2021). Anticipations in Macro: the Importance of Salience, Comprehensibility, and Actionability. *Keynote Lecture at RIDGE Forum on International Macro*.

Weber, M., B. Candia, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, O. Coibion, et al. (2023). Tell me Something I Don't Already Know: Learning in Low and High-Inflation Settings. *NBER, WP 31485*.

The Full, Persistent, and Symmetric Pass-Through of a Temporary VAT Cut

Online Appendix

Tiago Bernardino (IIES, Stockholm University)

Ricardo Duque Gabriel (Federal Reserve Board)

João Quelhas (Stockholm University)

Márcia Silva-Pereira (Nova SBE)

Appendix A Additional Figures

A.1 Background

Figure A.1: Headline and Food Inflation for Portugal

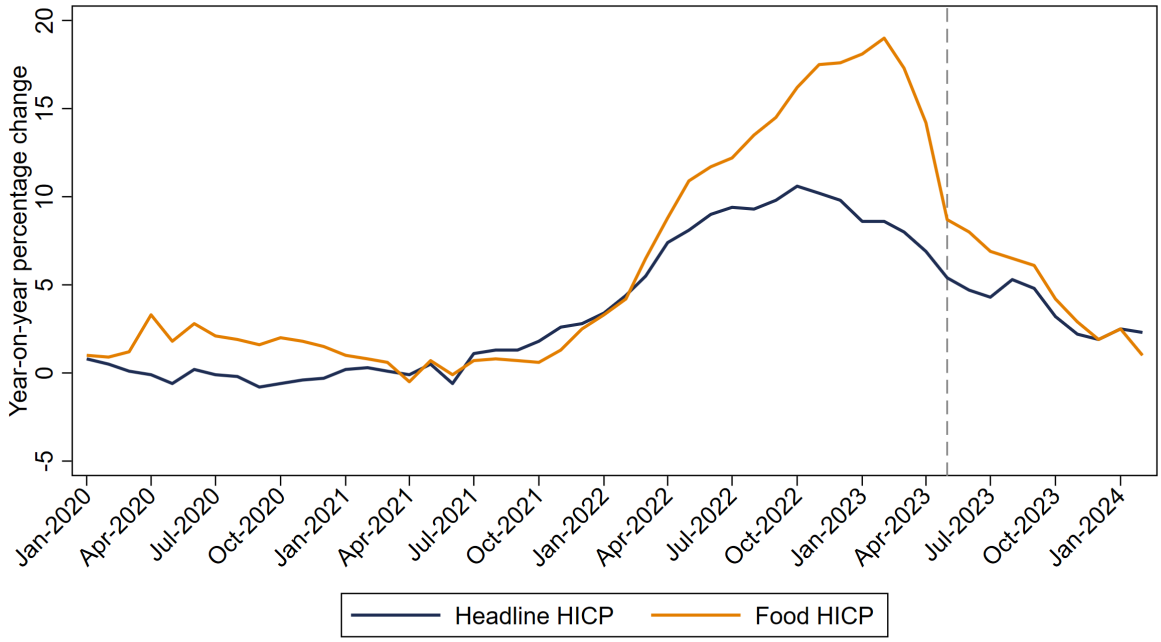

Notes: The figure shows the evolution of the year-on-year percentage change of the headline and the food Harmonized Index of Consumer Prices (HICP) baskets. The vertical gray dashed line corresponds to the implementation of the value-added tax cut policy in April 2023. *Source: BPstat, Banco de Portugal.*

Figure A.2: Media Coverage of the Temporary VAT Cut in Portugal

(A) Rejection by the Minister of Finance

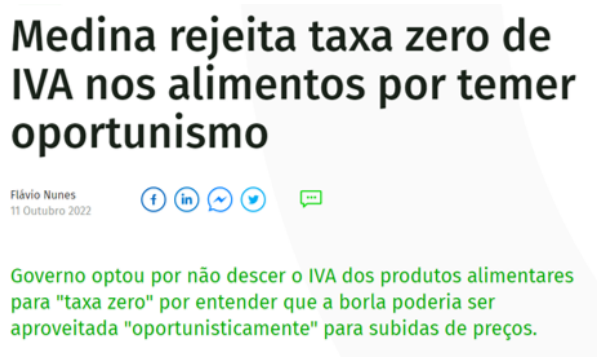

(B) Denial by the Minister of Finance

(C) Announcement by the Minister of Finance

Notes: Panel (A) refers to the first time the Minister of Finance rejected a VAT cut on food items. The title of the newspaper article is *"Medina [Minister of Finance] rejects a VAT cut on food items due to fears of opportunism."* Panel (B) refers to the second time the Minister of Finance denied that a VAT change was being prepared. The title of the newspaper article is *"Medina [Minister of Finance] insists that a VAT cut would not solve inflation in food products."* Panel (C) shows the moment when the VAT cut policy was announced, on March 24, 2023. The title of the newspaper article is *"Zero VAT. Medina [Minister of Finance] now believes in a measure that he criticized 10 days ago." Source: ECO, SIC and Renascença.*

Figure A.3: Salience of the Temporary VAT Cut in Portugal

(A) Prime Minister Visiting Supermarkets on the Day of the VAT Cut Implementation

(B) Labels in the Products Covered by the VAT Cut

(C) Advertisement for the Products Covered by the VAT Cut

Notes: Panel (A) shows the prime minister at the time, António Costa, visiting a supermarket and observing a display stating *"IVA 0%"* (in English, VAT 0%) next to zero VAT items. Panel (B) features a packaged fish with a sticker, stating *"IVA 0%"* (in English, VAT 0%). Panel (C) presents one advertisement made on a shelf displaying products covered by the VAT cut, saying *"Produtos com IVA 0%"* (in English, products with VAT 0%).

A.2 Descriptive Statistics

Figure A.4: Prices of Beans during the 2023 Temporary VAT Cut in Portugal

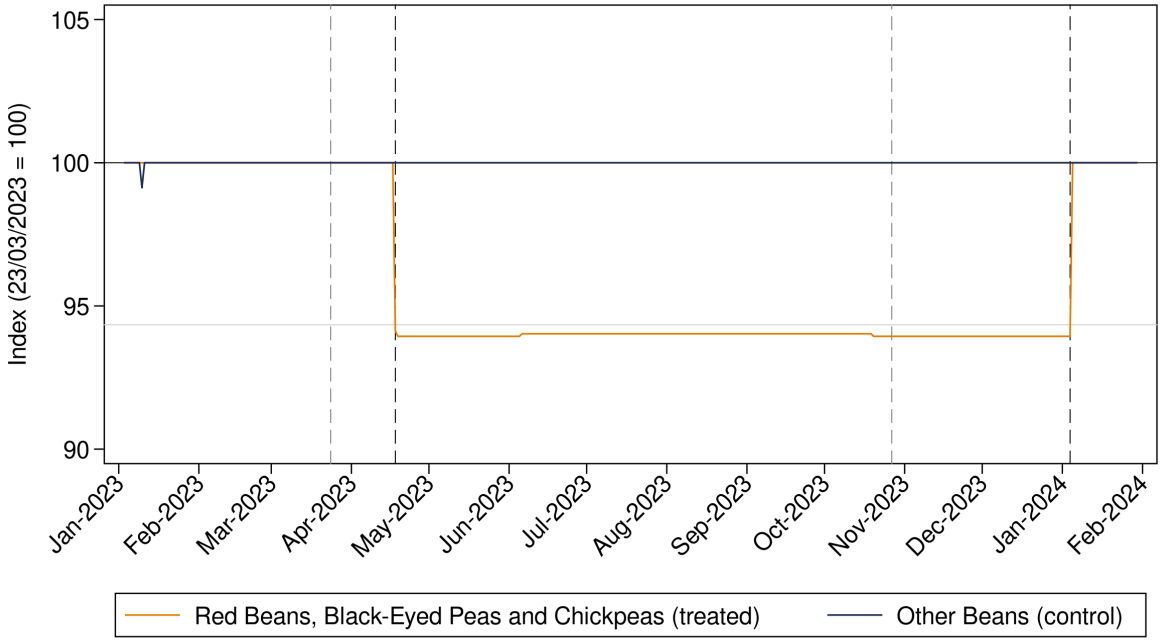

Notes: This figure shows the price dynamics of the beans that were covered by the VAT (e.g., red beans, black-eyed peas, and chickpeas; in orange) and the beans that were not (e.g., black beans and butter beans; in blue). The data represent the median value of the daily price indices for the items in each group from January 1, 2023, to January 31, 2024, normalized to 100 on the day before the unexpected announcement of the policy. The first dashed line indicates the announcement of the measure (March 24, 2023), the second dashed line marks the implementation day (April 18, 2023), the third dashed line indicates the official announcement of the reversal (October 27, 2023), and the fourth dashed line marks the reversal date (January 5, 2024). The horizontal solid line represents the price level corresponding to the full pass-through of the VAT cut policy. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

A.3 Full, Persistent, and Symmetric Pass-Through

Figure A.5: Pass-Through Estimates of the 2023 Temporary VAT Cut in Portugal

(A) Linear Trend Extrapolation

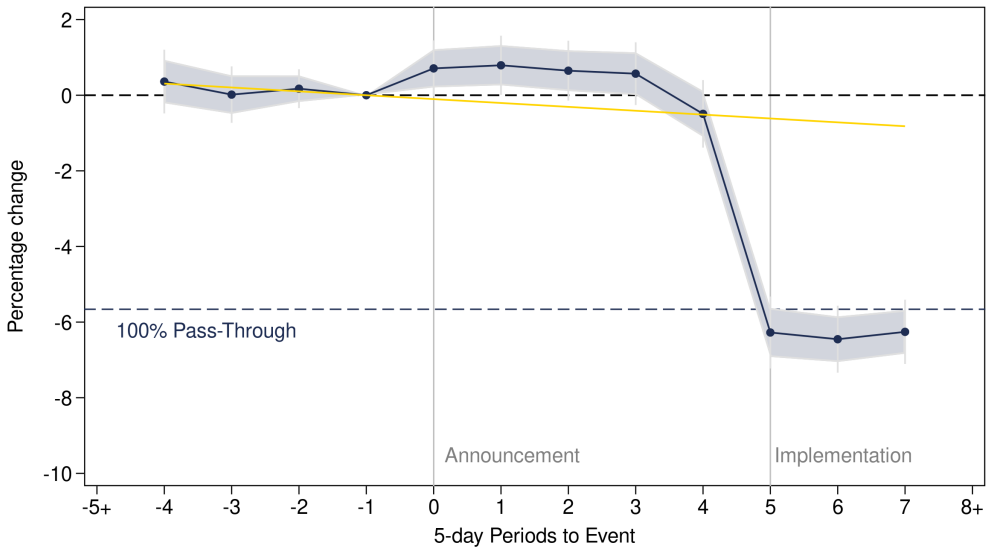

(B) Daily Data

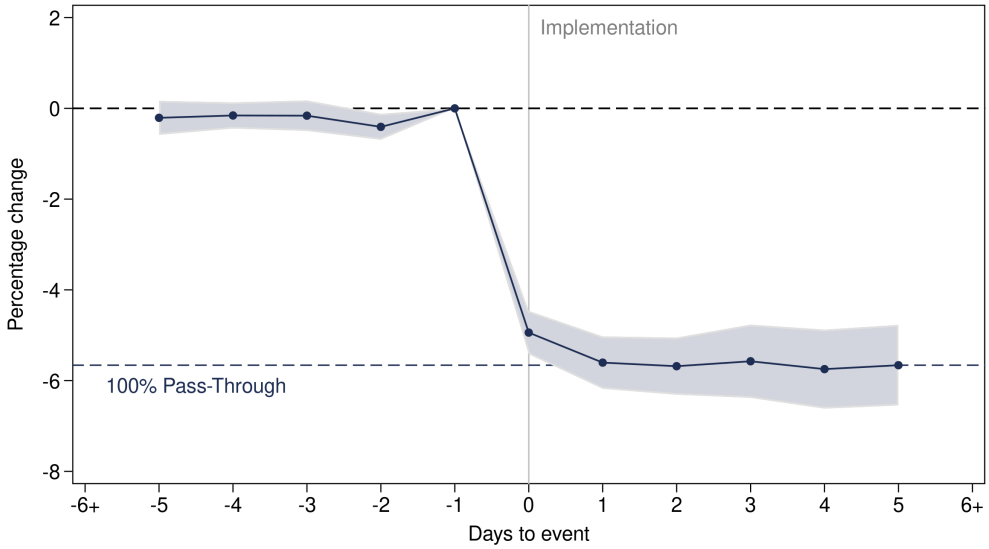

Notes: This figure shows the event-study estimates from Equation (1) after the policy announcement of the VAT cut. Panel (A) plots the coefficients and the 95% confidence interval without the trend correction, and the linear trend extrapolated, in yellow. The range of periods pre- and post-event considered in the analysis is M = 4 and G = 7. In this case, the p-value of the Wald test for "pre-trends" is 0.51, meaning that we do not reject the null hypothesis of "pre-trends." Panel (B) plots the coefficients and the 95% confidence interval using daily data. This means that in Equation (1), t corresponds to one day. The range of periods pre- and post-event considered in the analysis is M = 5 and G = 5. In this case, the estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). For both panels, item-level weights are used in the estimation, and the shaded areas correspond to 95% confidence intervals with standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

Figure A.6: Distribution of the Cumulative Percentage Change for Treated Items

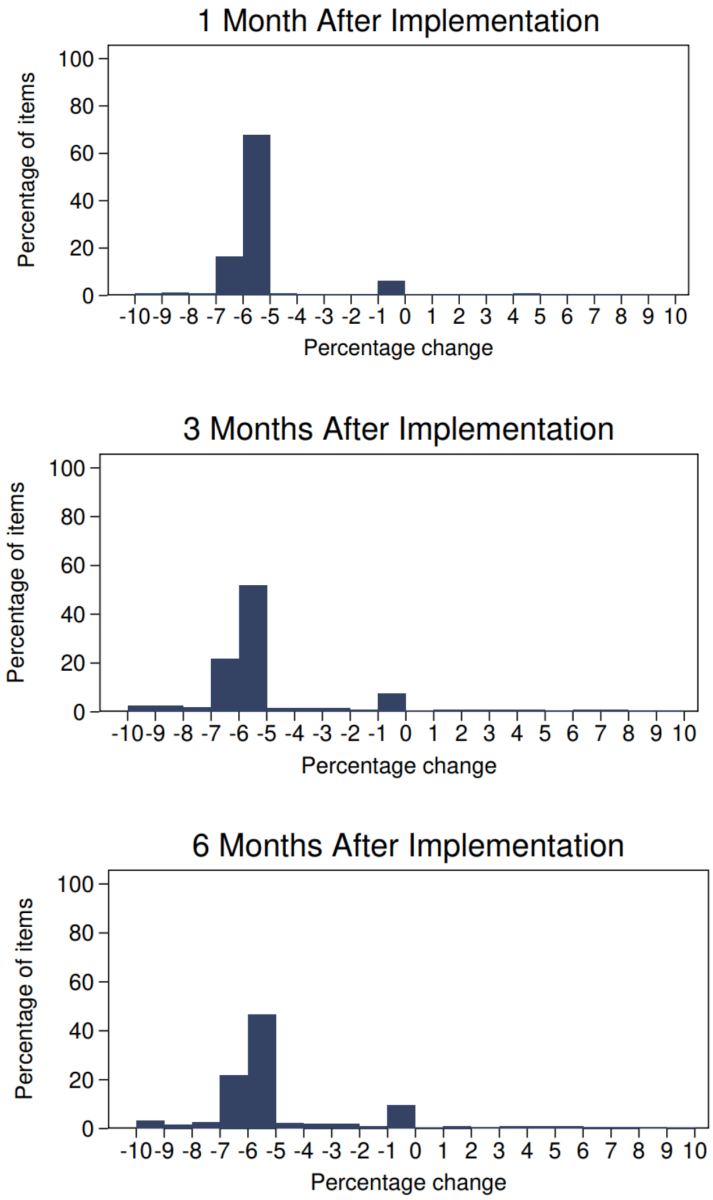

Notes: The Figure shows the magnitude of price changes comparing the price index level for each item on the day before the announcement of the policy and 1, 3, and 6 months after, respectively. It plots the distribution of the cumulative percentage change for treated items, with $p_{i,0}$ as the price before the policy's announcement. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

Figure A.7: Pass-Through Estimates of the 2023 Temporary VAT Cut Reversal in Portugal

(A) Linear Trend Extrapolation

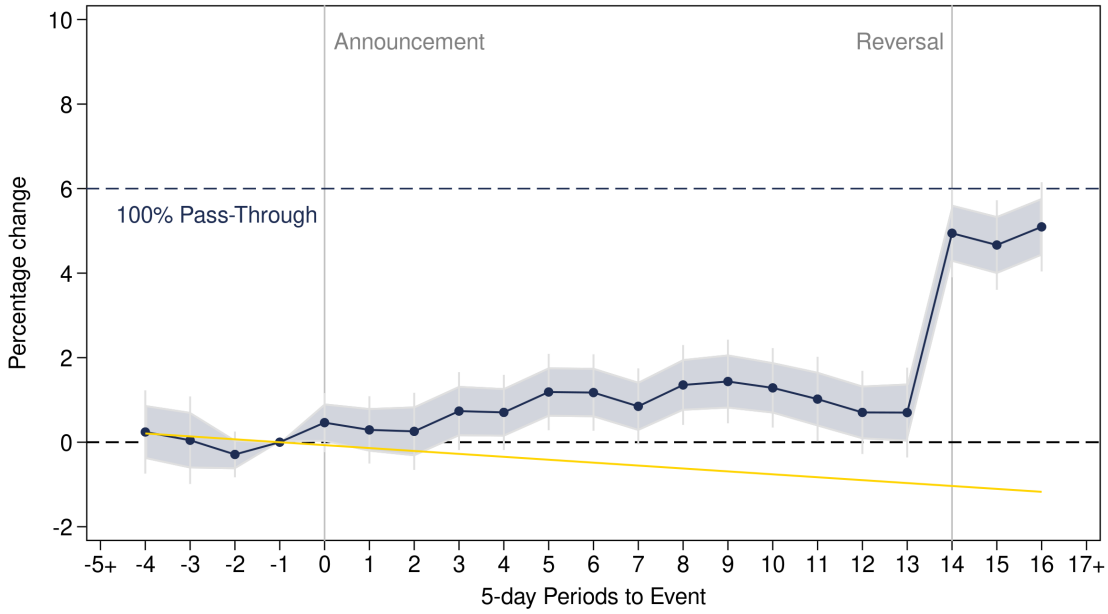

(B) Daily Data

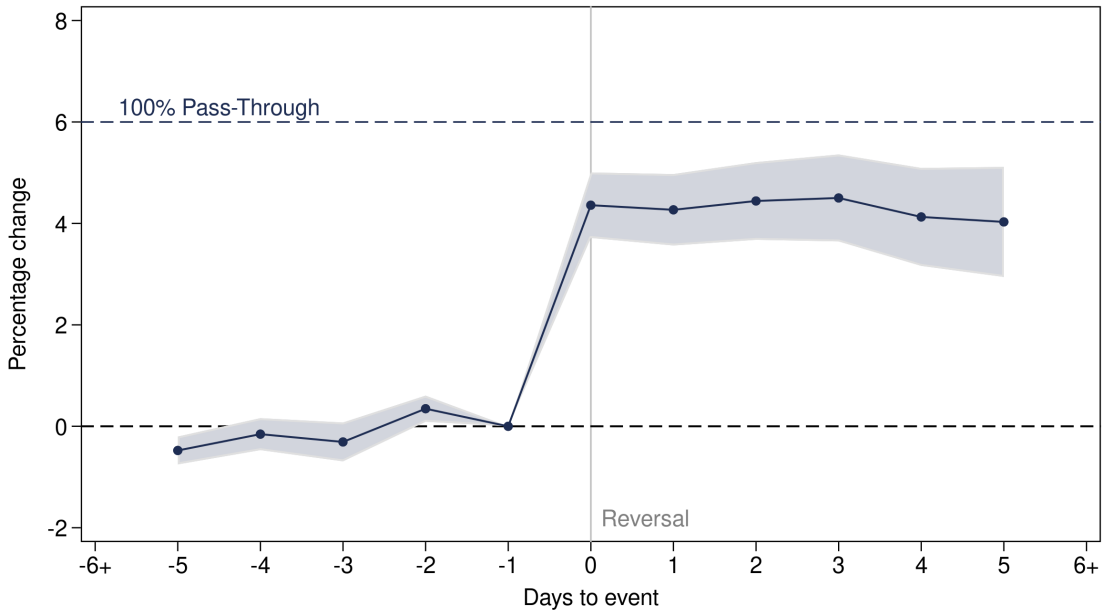

Notes: This figure shows the event-study estimates from Equation (1) after the reversal announcement. Panel (A) plots the coefficients and the 95% confidence interval without the trend correction, and the linear trend extrapolated, in yellow. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 16$. In this case, the p-value of the Wald test for "pre-trends" is 0.19, meaning that we do not reject the null hypothesis of "pre-trends." Panel (B) plots the coefficients and the 95% confidence interval using daily data. This means that in Equation (1), t corresponds to one day. The range of periods pre- and post-event considered in the analysis is $M = 5$ and $G = 5$. In this case, the estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). For both panels, item-level weights are used in the estimation, and the shaded areas correspond to 95% confidence intervals with standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

Figure A.8: Pass-Through Estimates of the 2023 Temporary VAT Cut in Portugal for Vegetable Oils

(A) Policy Implementation

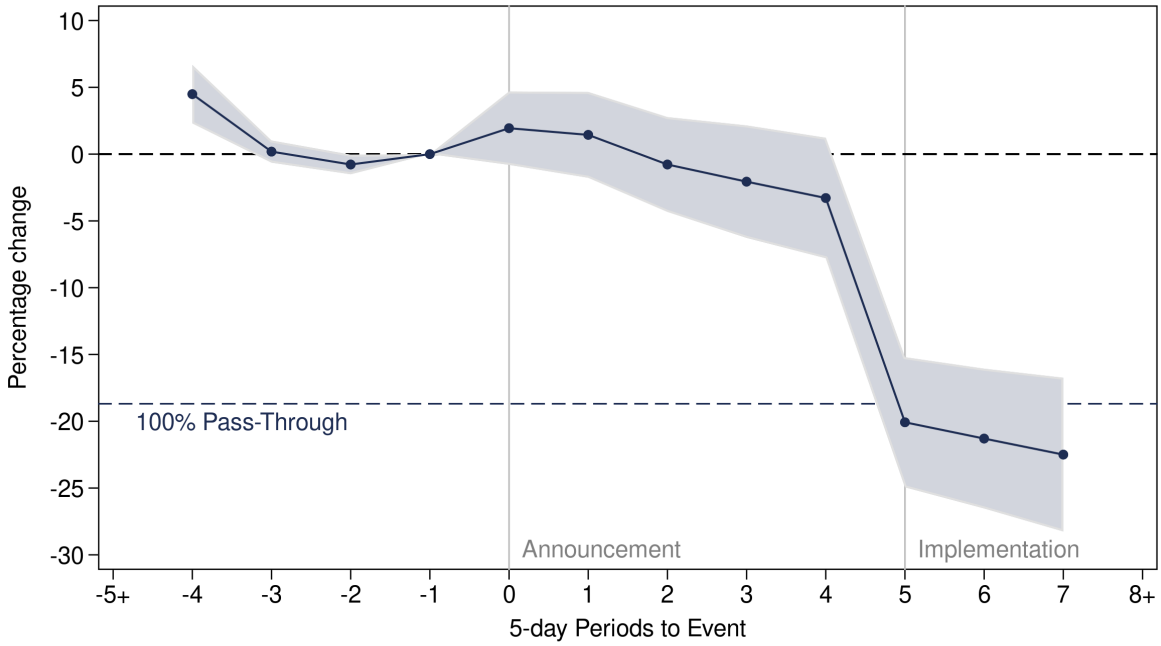

(B) Policy Reversal

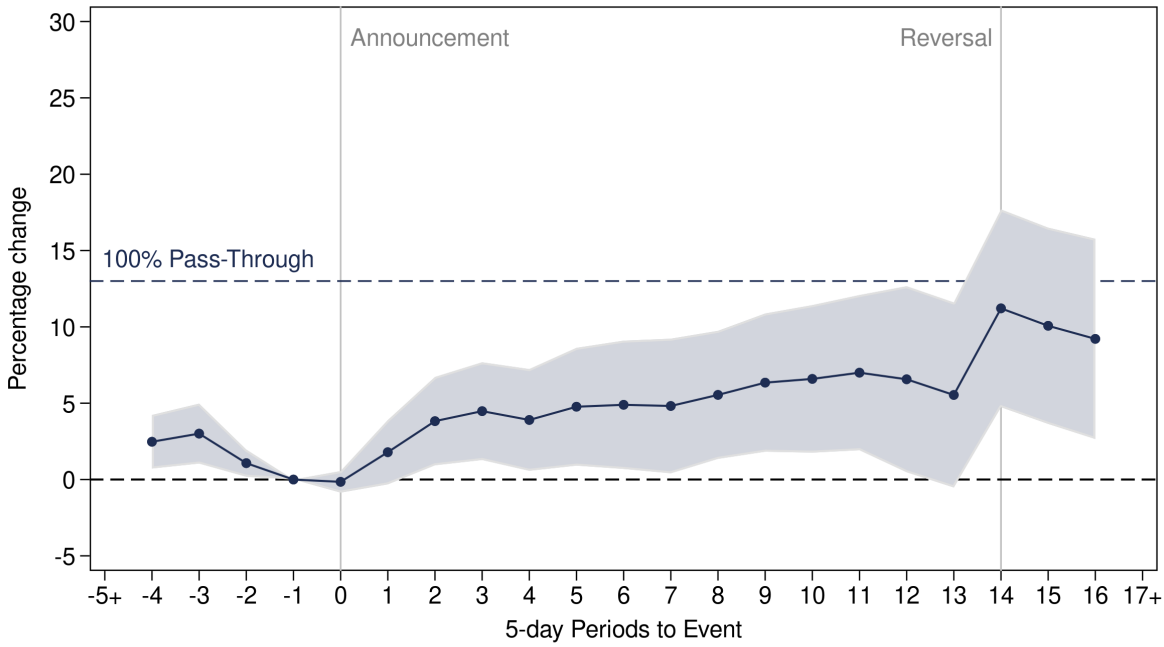

Notes: This figure replicates Figures 4 and 6 for the product category of vegetable oils. This category had a reduction in the VAT rate from 23% to 0% and then an increase up to 13% at the time of reversal. Panel (A) shows the event-study estimates from Equation (1) after the policy announcement. Panel (B) shows the event-study estimates from Equation (1) after the reversal announcement. The shaded areas correspond to 95% confidence intervals with the standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

A.4 Heterogeneity Analysis

Figure A.9: Oils and Fats

(A) Policy Reversal (with trend-correction)

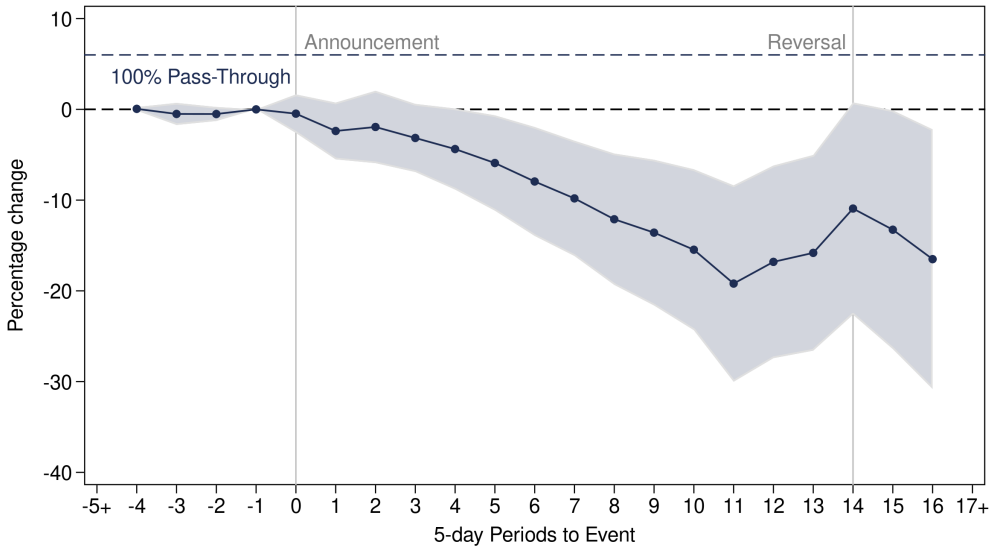

(B) Policy Reversal (without trend-correction)

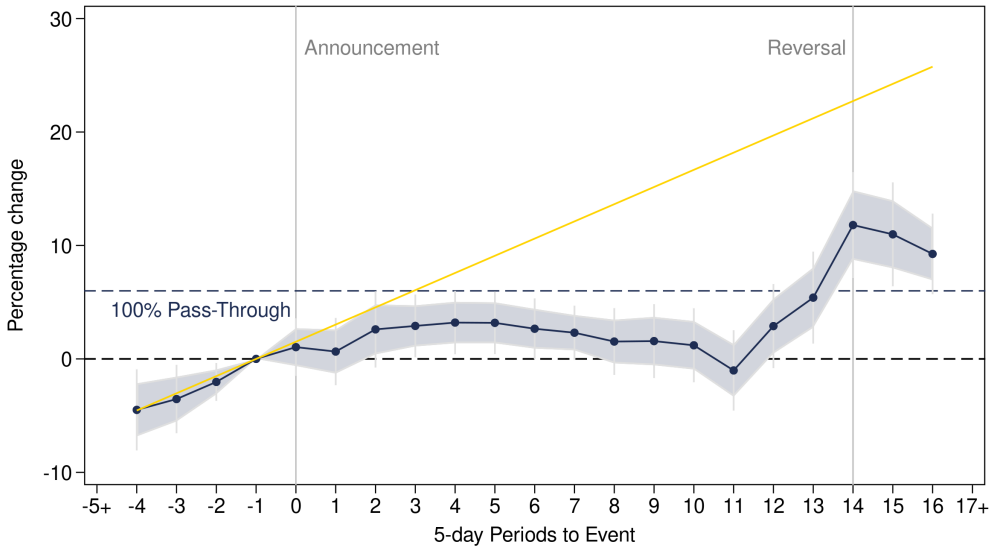

Notes: This figure replicates Figure 6 for the product category of oils and fats. Panel (A) shows the event-study estimates from Equation (1) after the reversal announcement. In this case, the estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). Panel (B) shows the event-study estimates without the trend correction, and the linear trend extrapolated, in yellow. The shaded areas correspond to 95% confidence intervals with the standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

Figure A.10: Heterogeneous Net Treatment Effects

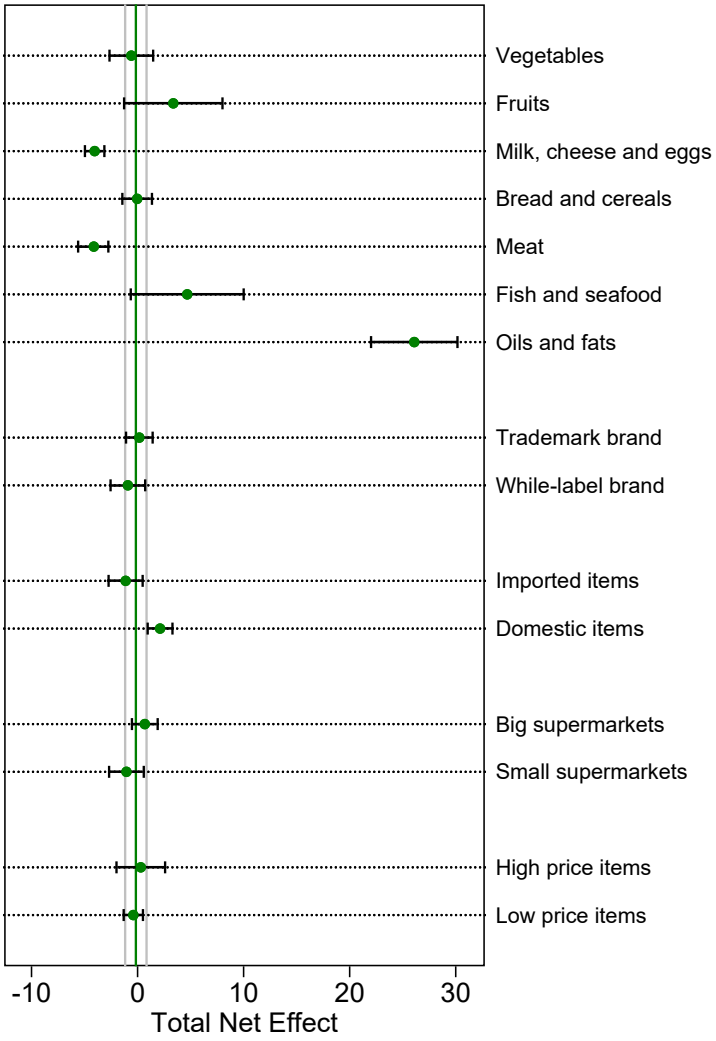

Notes: This figure shows the total net treatment effect along different product dimensions. The treatment corresponds to the event-study estimate from Equation (1) using subsamples created according to the different dimensions on the right: product categories (ECOICOP four-digit level), brand, origin, size of the supermarket, and price before the policy change. The vertical bar corresponds to the aggregate coefficient, along with the 95% confidence interval lower and upper bounds in gray. To estimate the total net effect, we use collapsed data at the monthly frequency, and the same trend-correction (Dobkin et al. 2018). The black solid lines correspond to the 95% confidence interval with standard errors clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

Figure A.10 shows the total net treatment effect between the policy announcement and its reversal for different product dimensions. For most ECOICOP categories and other product classifications, the total net treatment effect is not statistically different from the aggregate net treatment effect, represented by the vertical green line. However, some products deviate from this pattern – most notably, oils and fats, which exhibit a positive net total effect. As described in Section 5.4, and as Figure A.9 shows, before the announcement of the policy reversal, the price of this category was on an upward trajectory in comparison with the control group, due to a bad olives season harvest in 2023. This trend cannot be corrected in this exercise without further assumptions, implying that the net treatment effect estimated captures it and not only the VAT policy effect. In the right panel of Figure 7, the event-study regression is centered around the announcement of the policy reversal, which allows us to correct for the pre-trend before the announcement of the reversal using the Dobkin et al. (2018) correction.

Despite the persistency of our result and the fact that the net treatment effect is, on average, zero for several product categories and groupings, there are two caveats with this exercise. First, having a longer event study spanning over 8 months may introduce confounding effects that reduce its validity compared to looking only at the weeks immediately after each event. For example, there can be specific shocks to specific product categories that occurred between the policy announcement and its reversal that average out in the aggregate result, but lead to a net treatment effect different than zero when analyzing these heterogeneous net effects. Second, such a wide window of analysis, i.e., 8 months, leads to higher uncertainty in the estimation. As such, we use aggregated data at the monthly level instead of the 5-day periods. Without this, the confidence bands for the estimates would become excessively large due to the two-step estimation procedure that the correction uses.

A.5 Discussion

Figure A.11: Search Intensity on Google for "Value-Added Tax" and "VAT" in Portugal

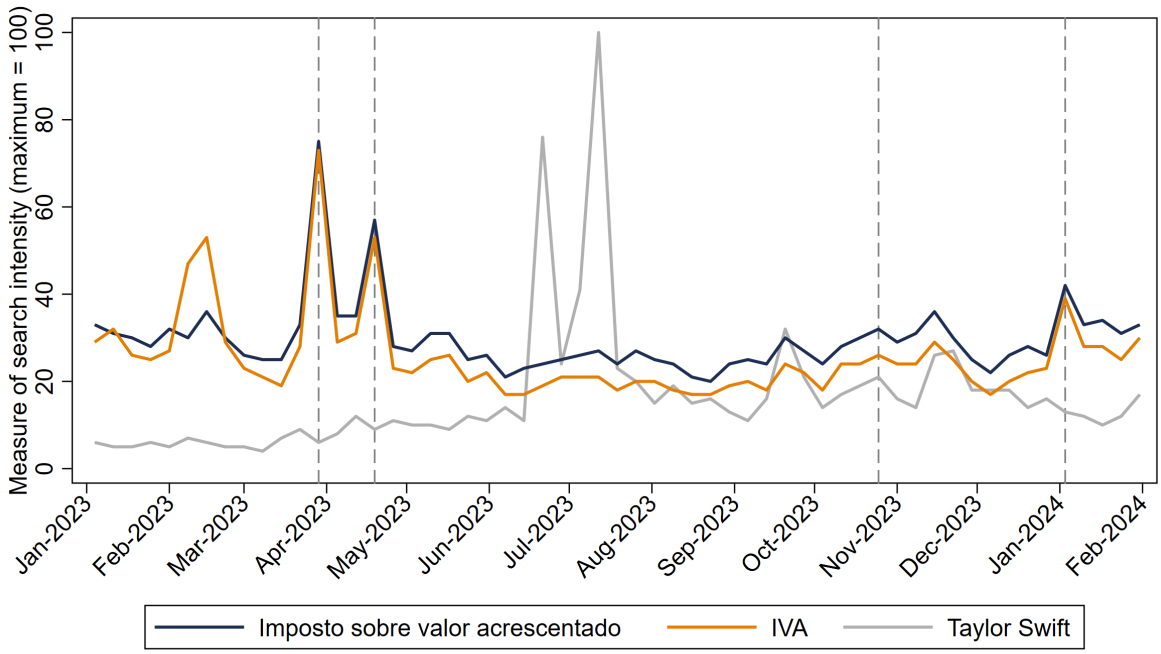

Notes: This figure shows the search intensity on Google for the expressions "Imposto sobre valor acrescentado" (in English, value-added tax) and "IVA" (in English, VAT). We also add the search intensity for "Taylor Swift" as a benchmark. The first gray dashed vertical line corresponds to the announcement of the VAT cut on March 24, 2023; the second gray dashed vertical line corresponds to the implementation date on April 18, 2023; the third gray dashed vertical line corresponds to the reversal announcement on October 27, 2023; and the fourth gray dashed vertical line corresponds to the reversal date on January 5, 2024. *Source: Google Trends.*

Figure A.12: Inflation Expectations over the Next 12 Months

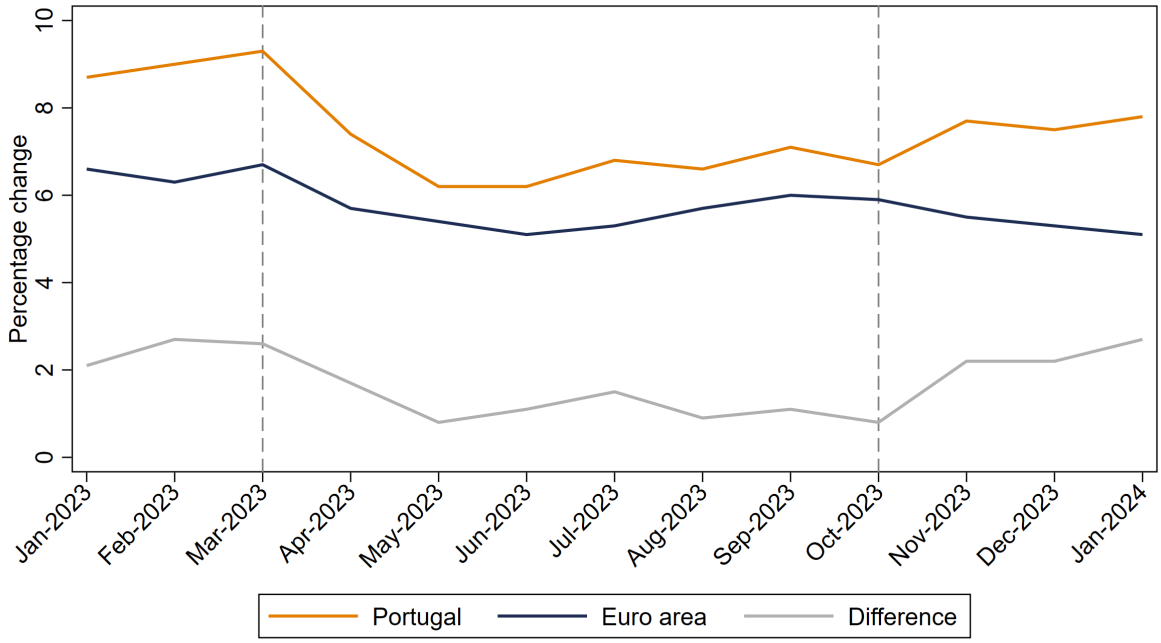

Notes: This figure shows inflation expectations over the next 12 months for Portugal and the euro area and the respective difference between the two series. We use weighted means, winsorized at the 2nd and 98th percentiles of each month and country. The euro area series includes the observations for 11 countries: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Portugal, Spain, and the Netherlands. *Source: European Central Bank's Consumer Expectations Survey.*

Figure A.13: Agricultural Producer Prices from Wholesale Markets

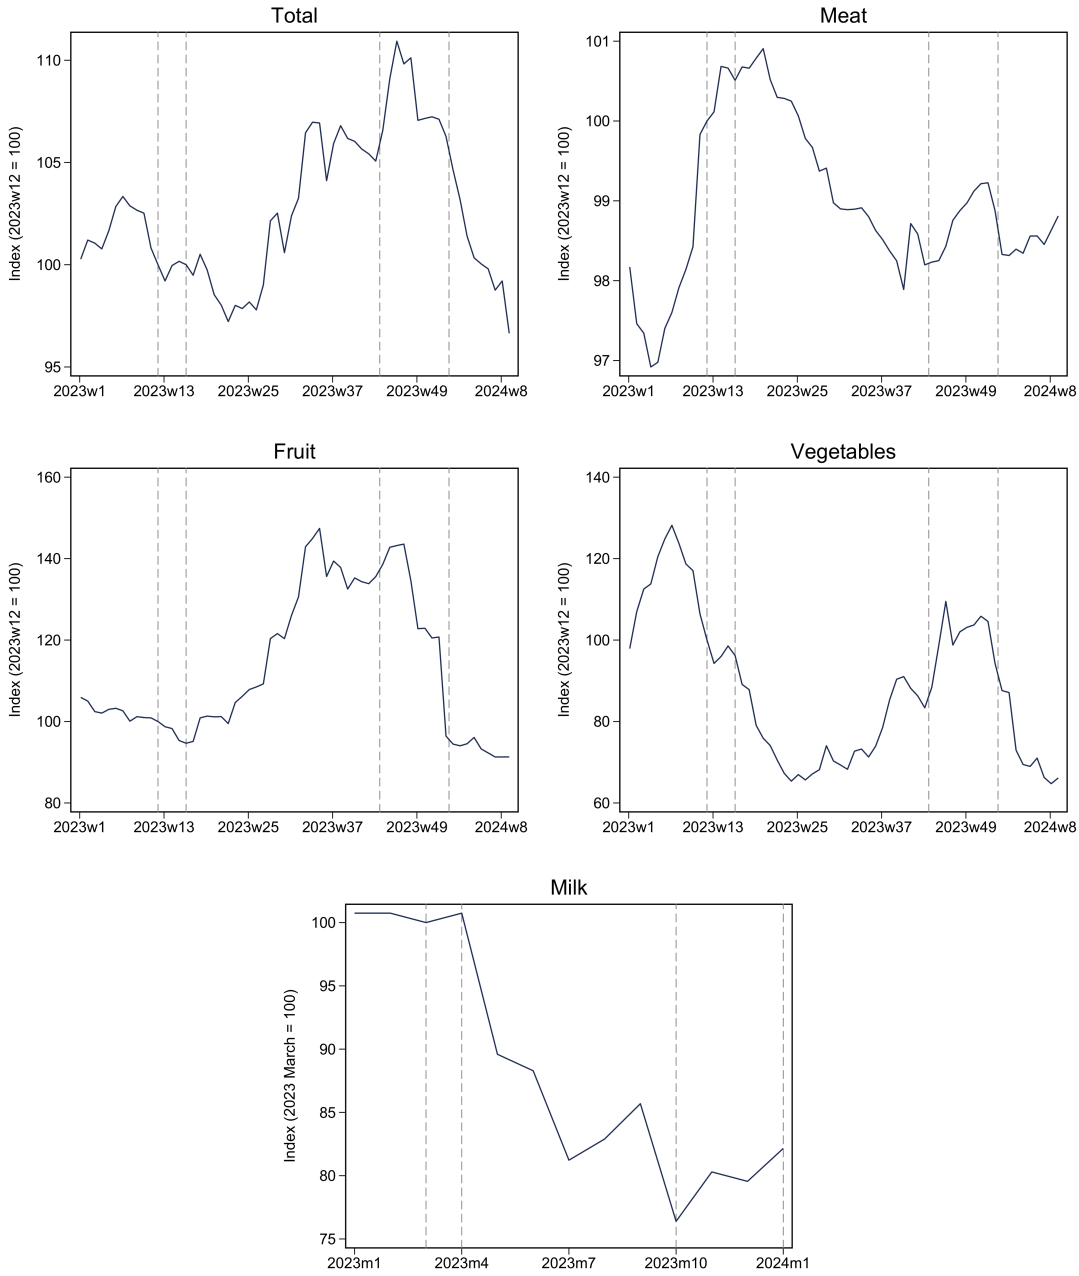

Notes: This figure shows the dynamics of agricultural producer prices in wholesale markets. The first four panels correspond to weekly data and have as a reference week the week preceding the policy announcement (week starting on March 20, 2024), and the last panel for milk corresponds to monthly data, indexed to March 2023. Figures for "Total" and "Meat" were produced by aggregating product-level data according to their 2023 ECOICOP five-digit level weights. The base period for the producer price indices is the week of the announcement—namely, week 12 of 2023. The four dashed lines correspond, respectively, to the week of the announcement, the week of implementation, the week of the announcement of the reversal of the policy until the end of the year, and the week of the reversal of the policy. Data description is available in Appendix B.6. *Source: Authors' calculations based on the price data of the Agricultural Markets Information System of the Portuguese Planning, Policy and General Administration Office.*

Figure A.14: Industrial Production Price Index for Manufacturing Food Industries

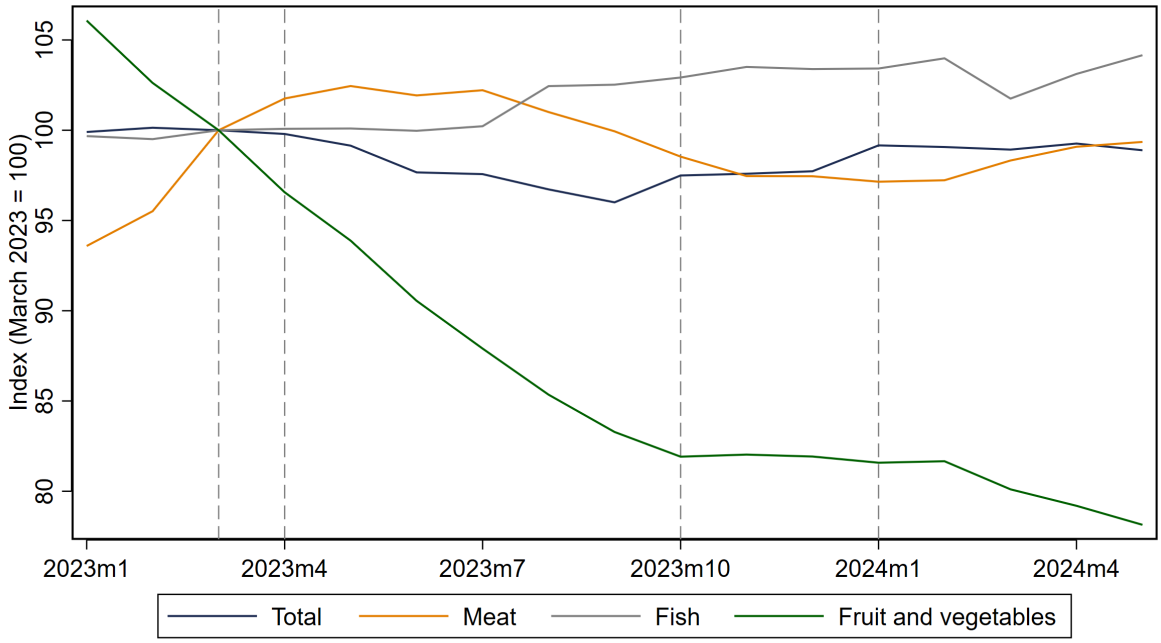

Notes: Data were retrieved from Statistics Portugal from January 2023 until May 2024, rebased to March 2023. The Industrial Production Price Index is available for all manufacturing industries, which include food industries. Meat, fish, and fruit and vegetables are the categories with the biggest weights in the representative consumer's consumption basket, as shown by the 2023 ECOICOP weights. *Source: Statistics Portugal.*

Appendix B Data

B.1 VAT Cut Basket

Table B.1 lists the names of the categories and respective products included in the VAT cut basket announced by the Portuguese government, as well as the size of the VAT rate change. The choice of the food products subject to the VAT cut was based on recommendations by the Directorate-General for Health, following the principles of a healthy and balanced diet and the most commonly consumed foods by the Portuguese population. All products were previously taxed at a 6% rate, with the exception of vegetable oils, and returned to the previous tax rate after the reversal. Vegetable oils had the standard tax rate of 23% before the VAT cut policy, and, at the time of the reversal, it increased to the intermediate tax rate of 13%. The analysis for this product category is performed separately and presented in Figure A.8.

Table B.1: List of Products Covered by the VAT Cut

Category	Products	Tax Rate Change
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%

Notes: This table shows the list of categories and respective products included in the VAT cut basket as well as the size of the VAT rate change. Vegetable oils were taxed at a 23% rate before the VAT cut and reverted to a 13% rate. All other categories had a change of 6% in both events. *Source: Law 17/2023, from April 14.*

B.2 Cleaning

Our data consist of web-scrapped prices from four 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 consist of 90,261 items (which correspond to product × supermarket—a product sold in different supermarkets has a different identifier) and 35,706,176 observations over time. We restrict to observations after January 1, 2023, leaving us with a total of 16,968,964 observations. During the first weeks of the policy, some goods raised doubts regarding their inclusion in the VAT cut basket. We exclude those items from our sample, decreasing the total number of items to 58,869. As Table B.1 indicates, the VAT cut size was 6% for all goods, except for vegetable oils, which were taxed at the rate of 23%. As such, we remove this product category from our sample to make our estimates consistent. This adjustment reduces the total number of items to 58,746. We also 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, excluding all drinks. This selection leaves us with a total of 43,283 products and 10,845,415 observations, corresponding to almost half of the total number of items in the raw data.

Table B.2 summarizes the above information.

Table B.2: Data-Cleaning Steps for the Baseline Specification

Step	Description	Items	Observations
0	Web scraped products	90,793	35,686,116
1	Exclude observations before January 1, 2023	63,218	16,966,732
2	Exclude treated products that raised doubts	59,430	15,976,570
3	Exclude treated products with a decrease of 23%	59,309	15,937,956
4	Exclude outlier observations	59,298	15,179,406
5	Exclude all drinks	43,283	10,845,415

Notes: This table shows the cleaning steps performed in the supermarket daily prices data set for the baseline specification present in Section 5. We define an item as a product × supermarket combination. Each observation corresponds to the posted price of one item in one day. Further details regarding the data set can be found in Banco de Portugal Microdata Research Laboratory (BPLIM) (2024). *Source: Author's calculations based on the BPLIM-SDP dataset.*

Table B.3 provides the number of items for each dimension considered in Section 5.4.

Table B.3: Number of Items for Each Category

	Control	Treated	Total
0111 Bread and cereals	**635**	**1,822**	**2,457**
01111 Rice	0	297	297
01113 Bread	0	752	752
01116 Pasta and couscous	635	773	1,408
0112 Meat	**70**	**747**	**817**
01121 Beef	0	148	148
01122 Pork	0	320	320
01124 Poultry	70	279	349
0113 Fish and seafood	**609**	**272**	**881**
01131 Fresh fish	265	58	323
01132 Frozen fish	344	214	558
0114 Milk, cheese and eggs	**938**	**5,321**	**6,259**
01141 Fresh whole milk	0	18	18
01142 Fresh skimmed milk	0	251	251
01144 Yogurt	0	2,042	2,042
01145 Cheese and curd cheese	0	2,231	2,231
01146 Other dairy products	938	696	1,634
01147 Eggs	0	83	83
0115 Oils and fats	**0**	**481**	**481**
01151 Butter	0	147	147
01153 Olive oil	0	334	334
0116 Fruit	**688**	**313**	**1,001**
01161 Fresh fruit	688	313	1,001
0117 Vegetables	**888**	**540**	**1,428**
01171 Fresh vegetables	626	376	1,002
01172 Frozen vegetables	140	81	221
01174 Potatoes	122	83	205
Other treated categories	6,834	684	7,518
Non-treated categories	22,441	0	22,441
Trademark Brand	27,051	7,810	34,861
White-Label Brand	6,052	2,370	8,422
Imported Items	16,961	4,826	21,787
Domestic Items	13,492	4,490	17,382
Non-classified Items	2,650	864	3,514
Big Supermarkets	18,334	6,046	24,380
Small Supermarkets	14,769	4,134	18,903
High-Price Items	15,406	4,788	20,194
Low-Price Items	17,697	5,392	23,089
Total	**33,103**	**10,180**	**43,283**

Notes: This table shows the number of treated and control items for each category across the dimensions that follow. The product category classification is based on the European Classification of Individual Consumption according to Purpose at the 4 and 5-digit level. Other treated categories are those with a small representativeness in the consumption basket and not considered for the heterogeneity exercise. These are other bakery and pastry products; other prepared and semi-prepared meat products; dried, salted, or smoked fish, and seafood; prepared and semi-prepared fish and seafood; other tubers and products derived; and dried vegetables. The non-treated categories are flours and other cereals; pizza and quiche; breakfast cereals; other cereal products; lamb and goat; other meats; edible offal; dried, salted, or smoked meat; fresh seafood; frozen seafood; other preserved or processed fish and seafood-based preparations; preserved milk; other milk products; margarine and other vegetable fats; other edible animal fats; frozen fruit; dried fruit and nuts; preserved fruit and fruit-based products; crisps; sugar; jams, marmalades, and honey; chocolate; confectionery products; edible ices and ice cream; artificial sugar substitutes; sauces and condiments; salt, spices and culinary herbs; baby food; ready-made meals; and other food products. Vegetable oils were excluded from the data set and analyzed separately. The white-label classification includes all items whose brand includes the name of the supermarket. The imported and domestic items are classified using the first three digits of the European Article Number (EAN) code, whenever available. Non-classified items correspond to those that do not have an EAN code available. We classify as big supermarkets those with a higher market share and as small supermarkets those with a smaller market share. High-priced and low-priced items are classified using as reference the mean price of the treated items before the policy announcement (3.8€). *Source: Author's calculations based on the BPLIM-SDP dataset.*

B.3 Weighting

Table B.4 reports the ECOICOP five-digit categories with at least one product in the VAT cut basket shown in Table B.1. It also presents the weights of the products included in the VAT cut basket within each category. For example, within the ECOICOP 01111 (Rice), all products had a decrease in the VAT rate from 6% to 0%. Within the ECOICOP 01116 (Pasta), only 78% of the products making up this category had a decrease in the VAT rate from 6% to 0%. These weights were provided by Statistics Portugal (Instituto Nacional de Esttística) and correspond to the ones used for the computation of the consumer price index (CPI) for Portugal in 2023.

To compute the weight of each item, we proceed as follows. We start by combining the weights in Table B.4 with the weight of each ECOICOP five-digit category in the food CPI basket to obtain the aggregate weight of the treated items and the control items for each category. Then, we divide the aggregate weight for each category × group by the total number of items, in each day. This ensures that we keep the structure of the food CPI in our analysis.

Table B.4: Weights for Each Product Category Included the VAT Cut Basket

ECOICOP 5-digit	Description	2023
01111	Rice	100%
01113	Bread	100%
01114	Other bakery and pastry products	8%
01116	Pasta and couscous	78%
01121	Beef	100%
01122	Pork	100%
01124	Poultry	97%
01128	Other prepared or semi-prepared meat products	68%
01131	Fresh fish	40%
01132	Frozen fish	77%
01135	Dried, salted or smoked fish, and seafood	100%
01136	Prepared and semi-prepared fish and seafood	62%
01141	Fresh whole milk	100%
01142	Fresh skimmed milk	100%
01144	Yogurt	100%
01145	Cheese and curd cheese	100%
01146	Other dairy products	26%
01147	Eggs	100%
01151	Butter	100%
01153	Olive oil	100%
01154	Vegetable oils	100%
01161	Fresh fruit	58%
01171	Fresh vegetables	69%
01172	Frozen vegetables	84%
01174	Potatoes	81%
01176	Other tubers and products derived	61%

Notes: This table reports the ECOICOP five-digit categories with at least one product in the VAT cut basket and the respective weight of these products with the category. *Source: Statistics Portugal.*

B.4 Five-Day Periods

Table B.5 presents the correspondence between the number of each five-day period and the days included in each one of them, as well as the main events related to the VAT cut policy.

Table B.5: Correspondence between Each 5-Day Period and the Starting and Ending Dates

Period	Starting Date	Ending Date	Period	Starting Date	Ending Date
1	Jan 3, 2023	Jan 7, 2023	40	Jul 17, 2023	Jul 21, 2023
2	Jan 8, 2023	Jan 12, 2023	41	Jul 22, 2023	Jul 26, 2023
3	Jan 13, 2023	Jan 17, 2023	42	Jul 27, 2023	Jul 31, 2023
4	Jan 18, 2023	Jan 22, 2023	43	Aug 1, 2023	Aug 5, 2023
5	Jan 23, 2023	Jan 27, 2023	44	Aug 6, 2023	Aug 10, 2023
6	Jan 28, 2023	Feb 1, 2023	45	Aug 11, 2023	Aug 15, 2023
7	Feb 2, 2023	Feb 6, 2023	46	Aug 16, 2023	Aug 20, 2023
8	Feb 7, 2023	Feb 11, 2023	47	Aug 21, 2023	Aug 25, 2023
9	Feb 12, 2023	Feb 16, 2023	48	Aug 26, 2023	Aug 30, 2023
10	Feb 17, 2023	Feb 21, 2023	49	Aug 31, 2023	Sep 4, 2023
11	Feb 22, 2023	Feb 26, 2023	50	Sep 5, 2023	Sep 9, 2023
12	Feb 27, 2023	Mar 3, 2023	51	Sep 10, 2023	Sep 14, 2023
13	Mar 4, 2023	Mar 8, 2023	52	Sep 15, 2023	Sep 19, 2023
14	Mar 9, 2023	Mar 13, 2023	53	Sep 20, 2023	Sep 24, 2023
15	Mar 14, 2023	Mar 18, 2023	54	Sep 25, 2023	Oct 1, 2023
16	Mar 19, 2023	Mar 23, 2023	55	Oct 2, 2023	Oct 6, 2023
17	Mar 24, 2023	Mar 28, 2023	56	Oct 7, 2023	Oct 11, 2023
18	Mar 29, 2023	Apr 2, 2023	57	Oct 12, 2023	Oct 16, 2023
19	Apr 3, 2023	Apr 7, 2023	58	Oct 17, 2023	Oct 21, 2023
20	Apr 8, 2023	Apr 12, 2023	59	Oct 22, 2023	Oct 26, 2023
21	Apr 13, 2023	Apr 17, 2023	60	Oct 27, 2023	Oct 31, 2023
22	Apr 18, 2023	Apr 22, 2023	61	Nov 1, 2023	Nov 5, 2023
23	Apr 23, 2023	Apr 27, 2023	62	Nov 6, 2023	Nov 10, 2023
24	Apr 28, 2023	May 2, 2023	63	Nov 11, 2023	Nov 15, 2023
25	May 3, 2023	May 7, 2023	64	Nov 16, 2023	Nov 20, 2023
26	May 8, 2023	May 12, 2023	65	Nov 21, 2023	Nov 25, 2023
27	May 13, 2023	May 17, 2023	66	Nov 26, 2023	Nov 30, 2023
28	May 18, 2023	May 22, 2023	67	Dec 1, 2023	Dec 5, 2023
29	May 23, 2023	May 27, 2023	68	Dec 6, 2023	Dec 10, 2023
30	May 28, 2023	Jun 1, 2023	69	Dec 11, 2023	Dec 15, 2023
31	Jun 2, 2023	Jun 6, 2023	70	Dec 16, 2023	Dec 20, 2023
32	Jun 7, 2023	Jun 11, 2023	71	Dec 21, 2023	Dec 25, 2023
33	Jun 12, 2023	Jun 16, 2023	72	Dec 26, 2023	Dec 30, 2023
34	Jun 17, 2023	Jun 21, 2023	73	Dec 31, 2023	Jan 4, 2024
35	Jun 22, 2023	Jun 26, 2023	74	Jan 5, 2024	Jan 9, 2024
36	Jun 27, 2023	Jul 1, 2023	75	Jan 10, 2024	Jan 14, 2024
37	Jul 2, 2023	Jul 6, 2023	76	Jan 15, 2024	Jan 19, 2024
38	Jul 7, 2023	Jul 11, 2023	77	Jan 20, 2024	Jan 24, 2024
39	Jul 12, 2023	Jul 16, 2023	78	Jan 25, 2024	Jan 29, 2024

Notes: This table shows the correspondence between each five-day period and the respective starting and ending dates. The policy announcement happens in period 17 (March 24, 2023—March 28, 2023). The policy implementation happens in period 22 (April 18, 2023—April 22, 2023). The reversal announcement happens in period 60 (October 27, 2023—October 31, 2023). The policy reversal happens in period 74 (January 5, 2024—January 9, 2024). Period 54 has six days.

B.5 Other Data Sources

HICP Aggregate Data. We collect aggregate HICP data for the treated products included in the VAT cut basket, using the ECOICOP five-digit classification, in Portugal and Spain. These data are sourced from Eurostat and collected by the national statistical offices that produce an index for all categories of ECOICOP for which the consumption is above 0.1% of the total household consumption expenditure in the country. The prices are collected every month and include products sold at online and physical stores.

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.[18] Data are available weekly for several product categories, species, regions, and local agricultural 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 five-digit ECOICOP category, which we also have in the SDP data set, and are presented in Table B.6.

We use data from January 2022 until February 2024. We keep products with observations available for at least two-thirds of the year, as seasonality may bias aggregate producer prices.

Additionally, we have monthly data from the same source for milk production by kilogram.

Table B.6: Selected Categories in Wholesale Markets

Category	ECOICOP	Weight (‰)	Number of items
Fresh or refrigerated fruit	CP01161	17.42	3,423
Poultry	CP01124	9.97	1,404
Fresh or refrigerated vegetables	CP01171	9.28	3,060
Beef	CP01121	8.26	12,217
Pork	CP01122	7.95	1,956
Olive oil	CP01153	3.24	410
Potatoes	CP01174	3.03	385
Eggs	CP01147	1.74	1,716

Notes: This table reports the selected categories of food products from the wholesale markets. Fresh or refrigerated vegetables exclude potatoes and other tubers. The third column represents the 2023 ECOICOP five-digit classification weight in permillage for each category in a representative consumption basket. The last column with the number of items corresponds to the number of product × market combinations for each category. *Source: Authors' calculations based on the price data of the Agricultural Markets Information System of the Portuguese Planning, Policy and General Administration Office.*

[18] Originally, in Portuguese, the *Sistema de Informação de Mercados Agrícolas* of the *Gabinete de Planeamento, Políticas e Administração Geral.*

Appendix C Robustness Checks

This section describes the several robustness checks that we perform. Appendix C.1 shows the robustness of the results of a different identification strategy. Appendix C.2 shows that our results are unchanged even when we use a SDiD method, a different estimation strategy that guarantees the parallel trends assumption. Appendix C.3 shows that the results are robust to the control group choice. Appendix C.4 shows that even with other cleanings in the data, the full pass-through is still there. And Appendix C.5 shows the results hold when using alternative price outcome variables in the event-study regression.

C.1 Identification Strategy: Portugal versus Spain

In this exercise, we use an alternative identification strategy. We use less granular data that prevent us from being as detailed as in the baseline exercise. However, we can still estimate the pass-through of the VAT cut policy when the policy is implemented and its persistency. We use the HICP data set provided by Eurostat. We take the price indices at the five-digit ECOICOP level of food categories for Portugal and Spain. These series have a monthly frequency.

Spain shares similar production and consumption structures with Portugal, and the two countries were affected similarly by recent economic shocks. Furthermore, the evolution of food prices in Portugal and Spain followed similar trends, making the parallel pre-trends assumption reasonable when using Spain as a control for Portugal.[19] Figure C.1 illustrates the price evolution of the food HICP in both countries.

Then, we apply a DiD identification, using the price dynamics in Spain as a control for the price dynamics in Portugal. Formally, we estimate the following equation:

$$P_{i,t} = \mu + \alpha C_i + \gamma T_t + \tau C_i \times T_t + \varepsilon_{i,t}, \tag{C.1}$$

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 1 for Portugal; and T_t is an indicator variable that takes the value of 1 after the VAT cut reduction. The treatment period is set to begin after May 2023, even though the policy was implemented on April 18, 2023. This is because the data used to compute the price indexes of April were collected in the first two weeks of the month and, thus, do not reflect the VAT cut, as one can see in Figure C.1.

[19] Spain adopted a similar policy of reducing the VAT on a narrower set of food items starting on January 1, 2023. Because of that, we use January 2023 as the first period in this exercise, when the trends are comparable.

Figure C.1: Price Evolution of the Basket of Treated Items in Portugal and Spain

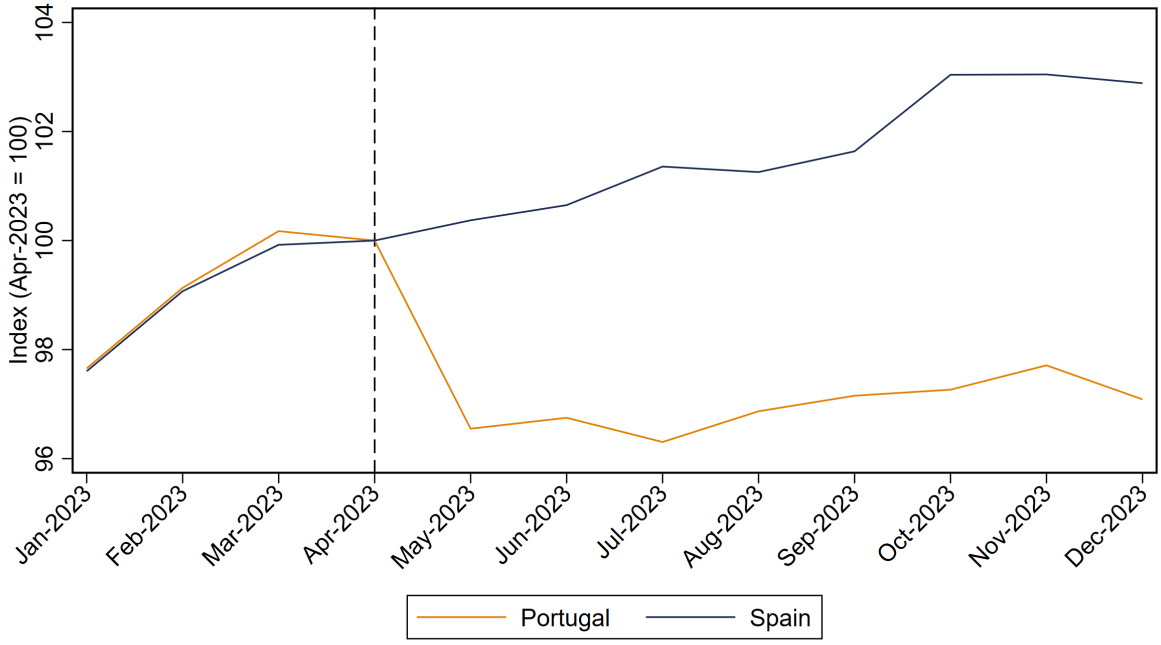

Notes: This figure shows the two series obtained by aggregating the price indices of the treated ECOICOP five-digit categories for Portugal (in orange) and Spain (in blue), using their corresponding weights for 2023, as defined by Eurostat. The gray dashed line corresponds to the implementation of the VAT cut in April 2023. *Source: Author's calculations based on the Eurostat HICP data.*

Table C.1 describes the estimation results of the above equation. In Column (1), we only use the country indicator in the regression. In Column (2), we only use the time indicator. In Column (3), we use both and their interaction at the same time. It also shows the robust standard errors, clustered at the ECOICOP five-digit level.

Table C.1: Robustness: Alternative Identification Strategy

	(1)	(2)	(3)
C_i	-1.88		1.77
	(2.111)		(2.163)
T_t		0.45	3.21**
		(1.240)	(1.193)
$C_i \times T_t$			-5.48***
			(0.953)
$N \times T$	1,188	1,188	1,188

Notes: This table shows the difference-in-differences estimation results for the alternative identification strategy. Robust standard errors are clustered at the ECOICOP five-digit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. *Source: Author's calculations based on the Eurostat HICP data.*

Our coefficient of interest is presented in Column (3). During the period when the policy was in effect, the CPIs of the affected product categories were, on average, 5.48% lower than

those of 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 5. 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.

C.2 Synthetic Difference-in-Differences

We employ the SDiD method to estimate the causal effects of the VAT cut in Portugal as a robustness check. This method combines elements of both Synthetic Control (SC) and DiD to provide a more flexible approach. By re-weighting and matching pre-exposure trends, SDiD reduces the reliance on the parallel trends assumption typical of DiD, and it remains valid even with additive unit-level shifts, which enhances the robustness for large-panel data.

In a conventional DiD analysis, we compare changes in outcomes over time between a treatment group and a control group, assuming that any differences are attributable to the first one. The SDiD offers a more adaptable approach: As a DiD, it accommodates situations where treated and control units have different pre-treatment trends but, like a SC, it optimally matches control units to treated units, significantly relaxing the need for parallel trends.

This dual approach allows SDiD to overcome common limitations of standard DiD and SC methods. For instance, standard DiD struggles to estimate causal effects if parallel trends do not hold, and SC requires that the treated unit falls within a "convex hull" of control units. The SDiD method, which we employ next, effectively addresses both of these issues.

We use the SDP data set with price data for both the treatment and control groups before and after the VAT cut. Then, we construct the SC group by employing statistical methods to match the pre-treatment characteristics of the treatment group. This step involves selecting and weighing appropriate control units with nonnegative weights to create a SC that closely resembles the treatment group. Finally, we estimate the causal effect by comparing the changes in the outcome variable over time between the treatment group and the SC

group after the intervention. The difference in these changes provides the estimated average treatment effect, which is comparable to the one obtained in Section 5.

We employ the SDiD estimator as developed by Arkhangelsky et al. (2021) and implemented for Stata by Clarke et al. (2023) using the *sdid* command. Our focus is on the case where there are multiple treatment units and a single treatment period.

Our analysis begins with a balanced panel comprising N units observed over T time periods. We use a consistent basket of products available throughout the entire analysis window. The outcome variable, the price index denoted P_{it}, is observed for each unit i in each period t. The binary treatment variable, denoted W_{it}, indicates whether unit i is treated at time t ($W_{it} = 1$) or not ($W_{it} = 0$). We assume a single adoption period for treated units, referred to by Arkhangelsky et al. (2021) as a "block treatment assignment." Once treated, units are assumed to remain exposed to the treatment indefinitely. The method requires at least two pre-treatment periods to determine appropriate control units (Clarke et al. 2023).

The objective of SDiD is to consistently estimate the causal effect of the VAT cut on the treated food items (W_{it}), yielding the average treatment effect on the treated (ATT). This approach is robust even if the parallel trends assumption does not hold between all treatment and control units. The estimation of the ATT proceeds as follows:

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

where the estimand is the ATT, derived from a two-way fixed-effects regression with optimally chosen weights, ω_i^{SDiD} and λ_t^{SDiD}. This approach allows for the inclusion of shared temporal aggregate factors through the estimation of time fixed effects γ_t and time-invariant unit-specific factors via unit fixed effects α_i. As is standard in fully saturated fixed-effects models, one α_i and one γ_t are normalized to zero to avoid multicollinearity. The presence of unit fixed effects ensures that SDiD matches treated and control units based on pre-treatment trends rather than levels, allowing for constant differences between these units.

The selection of unit weights, ω_i, ensures comparisons are made between treated units and controls that followed approximately parallel trends before the treatment. The selection of time weights, λ_t, prioritizes pre-treatment periods that are more similar to post-treatment periods, aiming for consistent differences between each control unit's post-treatment average

and pre-treatment weighted averages across all selected controls. Given the large number of treated and control units in our study, we confidently apply the estimator proposed by Arkhangelsky et al. (2021) and estimate its variance using a block bootstrap approach.

Figure C.2 illustrates the impact of the VAT cut on the price of food items using the SDiD method for the implementation and the reversal of the policy. The pre-treatment period demonstrates stable point estimates close to zero, indicating that there were no significant changes in the prices of the treated items relative to the SC group. This stability supports the validity of the parallel trends assumption, a crucial aspect of the SDiD methodology. Upon the implementation of the VAT cut on April 18, 2023, there is a sharp and significant decline in the price levels of the treated food items, relative to the control group, between -6.97% and -5.85%. This decline is aligned with the full pass-through of the VAT cut presented in our baseline exercise. Upon reversal, using as reference the respective preannouncement days, we find a similar behavior; an increase in the relative prices between 6.52% and 7.86%. Once again, the estimates are consistent with the baseline results.

Figure C.2: Robustness: Synthetic Difference-in-Differences

(A) Policy Implementation

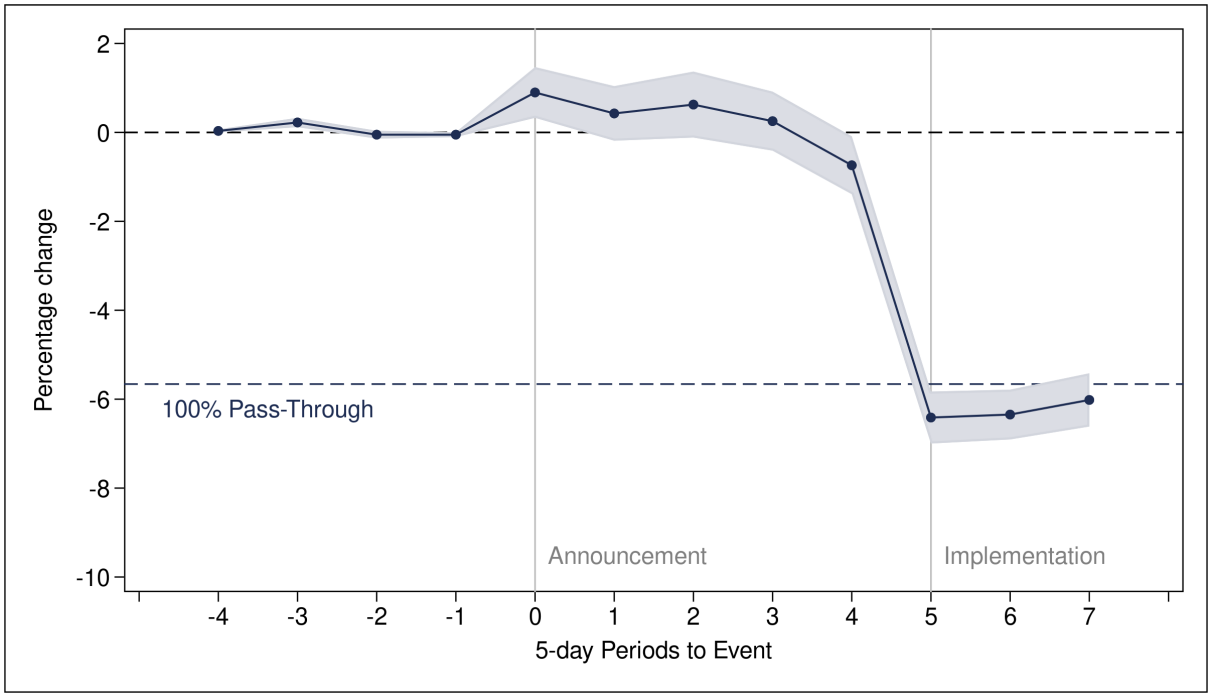

(B) Policy Reversal

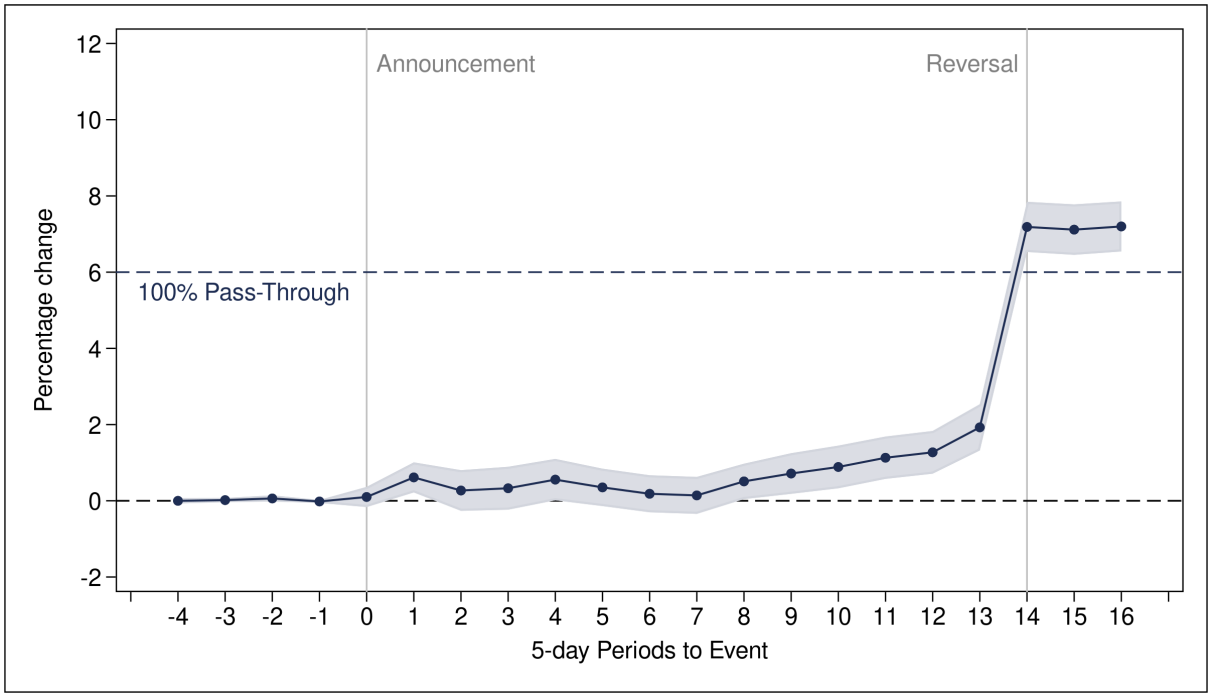

Notes: This figure shows the synthetic difference-in-differences estimates from Equation (C.2) in an event study-style output for the announcement (Panel A) and the reversal (Panel B) of the policy. In Panel (A), the balanced panel used for estimation comprises 7,634 units observed over 12 time periods. In Panel (B), the balanced panel used for estimation comprises 7,596 units observed over 17 time periods. The shaded areas correspond to 95% confidence intervals. Standard errors were computed using 1,000 bootstrap repetitions. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

C.3 Control Group

In our baseline estimation, we use all food products that were not included in the VAT cut basket as control, except for drinks. In this subsection, we show that our baseline result is robust to other control groups. We re-estimate Equation (1) using (1) all products sold by the supermarkets, (2) food and drinks, (3) nonfood products sold by supermarkets, and (4) only products that have a 6% VAT rate as control group. Table C.2 shows the point estimates and the standard errors clustered at the item level from the event study for the price change at the moment of the implementation and the reversal.

Table C.2: Robustness: Control Groups

Robustness	Description	Implementation	Reversal
1	All products as control	-5.873	6.032
		(0.451)	(0.837)
2	Only food products (incl. drinks)	-6.412	6.654
		(0.686)	(1.497)
3	Only non-food products	-5.271	6.529
		(0.459)	(0.858)
4	Only products with a 6% VAT rate	-5.016	4.524
		(0.722)	(1.603)

Notes: This table shows the event-study estimates from Equation (1) after the implementation and reversal, respectively, using different control groups. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 7$. The estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). Item-level weights are used. Standard errors in parentheses are clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

C.4 Data-Cleaning Methods

We also test different data-cleaning methods. The SDP data set has daily data on the prices of all products sold in a given supermarket. For different reasons, some goods do not have their prices available every day in the data. In the baseline exercise, we use the unbalanced panel including days when the prices of some goods are not available. In this Appendix, we repeat our exercise in the baseline using three alternative data-cleaning methods.

First, we impute prices such that we carry forward the information from the previous day for each product for a maximum of seven days if the price on the previous day with available information is equal to the price on the next day with available information until the product exits. Second, we also impute prices carrying them forward until the product exits without any time limitation for the missing gap. The third data treatment we test is imposing a balanced panel, where the basket of goods is constant. Table C.3 shows the point estimates and the standard errors clustered at the item level from the event-study for the price change at the moment of the implementation and the reversal.

Table C.3: Robustness: Data-Cleaning Methods

Robustness	Description	Implementation	Reversal
1	Data Imputation (1)	-6.236	6.049
		(0.702)	(1.646)
2	Data Imputation (2)	-4.520	4.736
		(0.592)	(1.543)
3	Balanced Panel	-5.780	7.062
		(1.952)	(3.392)

Notes: This table shows the event-study estimates from Equation (1) after the implementation and reversal, respectively, using different cleaning methods. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 7$. The estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). Item-level weights are used. Standard errors in parentheses are clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*

C.5 Outcome Variable

In the baseline specification, we use the price that consumers pay as the left-hand-side variable of the regression. In this Appendix, we repeat the exercise, using different outcome variables. First, we use the regular price, which is the price before sales and discounts. Second, we use the posted price per unit, to account for potential shrinkflation. Table C.4 shows the point estimates and the standard errors clustered at the item level from the event-study for the price change at the moment of the implementation and the reversal.

Table C.4: Robustness: Outcome Variables

Robustness	Description	Implementation	Reversal
1	Price per unit	-5.650	6.373
		(0.761)	(1.527)
2	Regular Price	-4.829	5.391
		(0.649)	(1.493)

Notes: This table shows the event-study estimates from Equation (1) after the implementation and reversal, respectively, using other outcome variables. The range of periods pre- and post-event considered in the analysis is $M = 4$ and $G = 7$. The estimated effect of the policy is the deviation from the extrapolated linear trend before the policy change, computed using four pre-event periods and a generalized method of moments estimator as in Dobkin et al. (2018). Item-level weights are used. Standard errors in parentheses are clustered at the item level. *Source: Authors' calculations based on the BPLIM-SDP dataset.*