

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

**Sticker Shocks: Using VAT Changes
to Estimate Upper-Level Elasticities of Substitution**

Bart Hobijn
Arizona State University

Fernanda Nechio
Federal Reserve Bank of San Francisco

October 2017

Working Paper 2015-17

<http://www.frbsf.org/economic-research/publications/working-papers/wp2015-17.pdf>

Suggested citation:

Hobijn, Bart, Fernanda Nechio. 2017. “Sticker Shocks: Using VAT Changes to Estimate Upper-Level Elasticities of Substitution.” Federal Reserve Bank of San Francisco Working Paper 2015-17. <http://www.frbsf.org/economic-research/publications/working-papers/wp2015-17.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Sticker Shocks: Using VAT Changes to Estimate Upper-Level Elasticities of Substitution*

Bart Hobijn

Fernanda Nechio

Arizona State University

FRB San Francisco

October, 2017

Abstract

We estimate the upper-level elasticity of substitution between goods and services of a nested aggregate CES preference specification. We show how this elasticity can be derived from the long-run response of the relative price of a good to a change in its VAT rate. We estimate this elasticity using new data on changes in VAT rates across 74 goods and services for 25 E.U. countries from 1996 through 2015. Depending on the level of aggregation, we find a VAT pass-through rate between 0.4 and 0.7. This implies an upper-level elasticity of 3, at the lowest level of aggregation with 74 categories, and 1 (Cobb-Douglas preferences) at a high level of aggregation that distinguishes 10 categories of goods and services.

JEL classification codes: E19, E21, D12.

Keywords: demand elasticities, multisector model, heterogeneity, aggregation, VAT rates.

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of San Francisco or the Federal Reserve System. This paper has benefited from comments and suggestions from Claudio Michelacci, three anonymous referees, Carlos Carvalho, Adam Shapiro, and seminar participants at HEC Montreal, UC San Diego, The Dutch Bank, EEA 2015, ES World Congress 2015, Federal Reserve Bank of San Francisco, SED 2015, and the Federal Reserve Macro System Meetings 2015. We would like to thank Eric Hsu and Ben Bradshaw for their excellent research assistance. E-mail: bhobijn@asu.edu, fernanda.nechio@sf.frb.org.

1 Introduction

Nested constant-elasticity of substitution (CES) preferences are the workhorse functional form used in multisector macroeconomic models.¹ CES preferences allow for a parsimonious representation of consumers’ willingness to substitute between varieties within a particular expenditure category, as well as their willingness to substitute across the broad classes of goods and services that make up expenditure categories.

The substitutability of varieties is determined by a lower-level elasticity of substitution while the substitutability across expenditure categories is parameterized by an upper-level elasticity of substitution. There are ample estimates of the elasticity of substitution between varieties of narrowly defined goods or services.² These are essentially estimates of lower-level elasticities of substitution and are conceptually different from the upper-level elasticity of substitution.³

There are, however, very few estimates of the upper-level elasticity of substitution. Those that exist are based on two different empirical approaches. The first approach is to estimate elasticities of substitution using micro-level data for many goods or a specific country.⁴ However, the expenditures covered in such analyses are still only a small subset of all goods and services over which aggregate preferences are commonly defined. The second approach is to estimate a classical demand system using macroeconomic data.⁵ Studies based on this approach run into the usual endogeneity problems that arise when movements in prices, which are used as explanatory variables, are jointly determined with variation in the dependent variables which are either quantities or expenditure shares.⁶

The literature reflects the absence of a broader-based estimate of the upper-level elasticity by resorting to analytically convenient calibrations for this parameter. For example, many studies use a unit elasticity. Others set it equal to the elasticity of substitution across varieties.⁷ These

¹Some examples that use such preferences include [Long and Plosser \(1983\)](#), [Aoki \(2001\)](#), [Woodford \(2003\)](#), [Carvalho \(2006\)](#), [Hobijn, Ravenna and Tambalotti \(2006\)](#), [Ngai and Pissarides \(2007\)](#), [Bouakez, Cardia and Ruge-Murcia \(2009\)](#), [Nakamura and Steinsson \(2010\)](#), [Carvalho and Nechio \(2011\)](#), [Eusepi, Hobijn and Tambalotti \(2011\)](#), [Midrigan \(2011\)](#), and [Carvalho and Nechio \(2016\)](#).

²See, for example, [Nevo \(2001\)](#) for the market for cereals, [Manning et al. \(1987\)](#) for the demand for medical services, [Petrin \(2002\)](#) for the market for cars and minivans, [Leslie \(2004\)](#) for Broadway plays, and [Broda and Weinstein \(2006\)](#), [Simonovska and Waugh \(2014a\)](#), [Simonovska and Waugh \(2014b\)](#), [Feenstra et al. \(2014\)](#), and [Imbs and Mejean \(2015\)](#) for tradable goods.

³In particular, if one assumes that price movements in these narrow goods categories do not significantly affect the overall price level ([Dixit and Stiglitz, 1983](#)).

⁴See, for example, [Broda and Weinstein \(2006\)](#), [Gabriel and Reiff \(2010\)](#), and [Feenstra et al. \(2014\)](#).

⁵For example, the “Almost Ideal Demand System” introduced in [Deaton and Muellbauer \(1980\)](#).

⁶See [Berry \(1994\)](#) and [Berry, Levinsohn and Pakes \(2004\)](#) for a thorough discussion of simultaneity bias in such models.

⁷Examples of papers that have followed one or the other calibration approach can be found in the macroeconomics New Keynesian literature, such as [Carvalho \(2006\)](#) and [Carvalho and Nechio \(2011\)](#), among others, and in the trade literature (see [Costinot and Rodríguez-Clare, 2014](#) for a summary).

two approaches have resulted in calibrations of the upper-level elasticity ranging from 1 to as large as 11. In addition, such calibrations are typically done with little discussion about the level of aggregation of the expenditure categories in the analysis.

In this paper we provide an estimate of the upper-level elasticity of substitution that is different from the two aforementioned approaches for two reasons. First of all, it is comprehensive in its scope. It is based on data that cover *all* expenditures in the basket of goods and services that make up the Harmonised Index of Consumer Prices (HICP), which is the main inflation gauge in the European Union. Secondly, it does not suffer from endogeneity bias. This is because we identify the upper-level elasticity of substitution by estimating the *long-run* response of relative prices to changes in relative Value Added Tax (VAT) rates. These VAT changes are the proverbial “sticker shocks” referred to in the title.

Other papers in the literature have studied the impact of VAT changes on inflation.⁸ Contrary to most other studies, however, our analysis covers *all* goods and services in the HICP as well as a large sample of countries.⁹ More importantly, our objective is to use VAT changes for the estimation of the upper-level elasticity of substitution for aggregate nested CES preferences.¹⁰

The main intuition behind our method is that an increase in the VAT is equivalent to an exogenous rise in the marginal cost of production. This induces monopolistically competitive firms to raise their prices. This price increase, in turn, results in a decline in demand, the size of which depends on the elasticity of substitution that we are interested in.

Based on this intuition, we derive the *long-run* response of relative prices to changes in relative VAT rates in a model of price setting under monopolistic competition, where demand is determined by two-tier nested CES preferences and goods are produced using a Cobb-Douglas production technology. This framework is very general and nests those commonly applied in Real Business Cycle (RBC) models, New-Keynesian models, and trade models.¹¹ Thus, our estimate is applicable under a very general set of assumptions that covers many macroeconomic models.

⁸Studies that also do so are, for example, [Karadi and Reiff \(2007\)](#), [Carare and Danninger \(2008\)](#), [Gabriel and Reiff \(2010\)](#), [Gautier and Lalliard \(2014\)](#), [Benedek, de Mooij and Wingender \(2015\)](#). Similarly, papers in the trade literature have used changes in tariffs to identify trade elasticities.

⁹The countries in our sample are, by construction, those that have VAT rates. In such countries national accounts are generally constructed from the production side of the accounts and detailed expenditure data is only collected at a low frequency. Even in countries that do have high-frequency expenditure data, the fluctuations in these data at low levels of aggregation are imprecisely measured ([Wilcox, 1992](#)).

¹⁰Akin to our analysis, in an event study, [Cashin and Unayama \(2016\)](#) use a pre-announced VAT rate change in Japan to infer on Japanese consumers’ intertemporal elasticity of substitution.

¹¹Three examples of RBC, New Keynesian, and trade models that fit into the framework we use are those in [Long and Plosser \(1983\)](#), [Aoki \(2001\)](#), and [Costinot and Rodríguez-Clare \(2014\)](#), respectively.

In our model, firms face decreasing returns to scale with respect to factors that are adjustable in the short-run, which is labor in our case. Because of this, the decline in demand leads to a reduction in marginal costs, which partly offsets the price increase. As a result, our model predicts an incomplete pass-through of VAT changes due to the interaction of the elasticity of substitution and the returns to scale of the production technology.

We show that, within this framework, the reduced-form long-run response of relative prices to changes in relative VAT rates depends on two parameters; (i) the labor elasticity of output and (ii) the upper-level elasticity of substitution that is the focus of our analysis. Hence, for a given labor elasticity of output, we can back out the upper-level elasticity of substitution from the long-run effect of VAT changes on inflation.

In order to estimate this reduced-form long-run response, we use the local projection method introduced in [Jordà \(2005\)](#). This method estimates the dynamic response of a variable (here relative prices) with respect to an exogenous movement in another variable (here VAT rates) over different horizons using a simple least-squares regression framework. The regression specification is adjusted based on the horizon over which the response is estimated.

In the context of our problem, this local projection method boils down to running a simple panel-data regression of goods-specific changes in VAT rates on cumulative goods-specific inflation rates for different horizons. Based on these regression results, we use the long-term effects of VAT rate changes on inflation to back out the upper-level elasticity of substitution. In principle, this panel-data model can be estimated on a country-by-country basis. However, since VAT changes occur rather infrequently, this results in imprecise estimates of the reduced-form elasticity. For this reason we pool our regression across countries.¹²

The estimation of our panel-data model requires disaggregated data on inflation and VAT changes by expenditure category for each country. For this purpose, we construct a data set in which we measure inflation using monthly inflation rates for all categories of expenditures that are included in the HICP. Because monthly data on VAT rates by expenditure category and country are not readily available, we obtain them from two administrative sources; [European Commission \(2015\)](#) and [Eurostat \(2015\)](#). The result is a dataset of monthly inflation and VAT rates for 74 expenditure categories and 25 countries, that covers the period from January 1996 to January 2015.¹³

In order to provide useful estimates of the elasticity of substitution that can be used for the calibration of preference parameters at different levels of aggregation, we estimate the upper-

¹²For the estimates obtained using the local projection method to be consistent, the changes in VAT rates need to be exogenous, uncorrelated with future VAT changes, and expected to be permanent. Throughout the text, we present evidence that this is indeed the case.

¹³A similar data set was concurrently constructed by [Benedek, de Mooij and Wingender \(2015\)](#).

level elasticity of substitution (separately) for the three levels of aggregation at which the data are provided. These correspond to 1-, 2-, and 3-digit expenditure classifications which are the basis for the HICP measure.

Depending on the level of aggregation, we find a VAT-pass-through rate between 0.4 and 0.7. This implies a point estimate at the lowest level of aggregation is of an elasticity of substitution approximately equal to 3. Given the standard error around this estimate, the upper bound of the 95% confidence interval on the latter estimate is about 5. Our estimate of the upper-level elasticity of substitution at the highest level of aggregation, at which expenditures are split up into 10 categories, is one. This is in line with the use of Cobb-Douglas preferences at this high level of aggregation.

These results provide useful guidance for the choice of the value of the upper-level elasticity of substitution in multisector macroeconomic models. The properties of such models tend to vary substantially with the value of this elasticity. For example, it is important in New-Keynesian models with sticky prices, in which it influences the size of the distortion in relative prices due to nominal rigidities. This is the distortion that monetary policy (partially) offsets in these models.¹⁴ It also affects the magnitude of the gains from trade.¹⁵

The bottom line is that, depending on the level of aggregation at which the upper-level expenditures are defined, a reasonable choice for the upper-level elasticity of substitution is between 1 and 3. For any choice higher than 5 there is little support in the data.

2 Model

We consider a partial equilibrium model of price setting where demand functions are determined by two-tier nested CES preferences and production is done using a Cobb-Douglas technology. This setup nests both price setting under monopolistic competition with sticky prices, as discussed in [Woodford \(2003\)](#), as well as (in the limit) price setting under perfect competition, as for example in [Long and Plosser \(1983\)](#). We follow [Karadi and Reiff \(2007\)](#) and add a VAT rate that affects the firms' price setting decisions.

Given this setup, we take the following approach. We first derive the price setting decisions of firms and solve for the goods' relative price that will prevail in steady state. We then show that, for each good, the response of its relative price with respect to its VAT rate only depends on two parameters. The first is the curvature of the production function with respect to flexible

¹⁴See, for example, [Blanchard and Galí \(2007\)](#) and [Carvalho and Nechio \(2016\)](#).

¹⁵See [Costinot and Rodríguez-Clare \(2014\)](#) for a discussion of the effect of the upper-level elasticity of substitution on the gains from trade.

inputs, which in our model is pinned down by the labor share. The second is the elasticity of substitution between goods, which is our parameter of interest.

2.1 Price setting with nested CES preferences and Cobb-Douglas technology

The economy that we consider is one in which consumers derive utility from the consumption of different types of goods, indexed by $j = 1 \dots J$. A continuum (of measure one) of varieties of each of these goods is supplied. Each variety $i \in (0, 1)$ is highly substitutable for the others within the goods category. These varieties constitute the lower level of the nested CES preferences that we analyze. The goods represent the higher level of these preferences. Our parameter of interest is the elasticity of substitution between goods at this higher level of CES preferences.

2.1.1 Price setting decisions by producer of variety i

Each firm produces a variety i of good j in period t , Y_{ijt} , using a decreasing returns to scale Cobb-Douglas production technology. Labor, which is the sole input in production, is mobile across firms that produce different goods and varieties. The amount of labor used in the production of variety i of good j at time t is L_{ijt} . At a given total factor productivity level, A_t , output of the firm equals

$$Y_{ijt} = A_t L_{ijt}^{1-\alpha}. \quad (1)$$

This firm hires labor at the nominal wage rate, W_t , which it takes as given. Consequently, the marginal cost of producing an extra unit of output for the firm is

$$MC_{ijt} = \frac{1}{1-\alpha} W_t \left(Y_{ijt}^{\frac{\alpha}{1-\alpha}} / A_t^{\frac{1}{1-\alpha}} \right). \quad (2)$$

Since the varieties are close, but not necessarily perfect, substitutes, the firm is a monopolistic competitor. This means that this firm is not a price taker but, instead, chooses a point on its variety-specific demand curve.

For our representation of this demand curve, we denote the retail price that the firm charges for its variety, including the VAT, by P_{ijt} and the price of good j across all varieties by P_{jt} .¹⁶

¹⁶In the European countries in our sample, prices are quoted including VAT charges, as opposed to the United States, where most price quotes exclude sales tax. Therefore, P_{ijt} is the price that affects the household's cost of living, and determines the consumer's demand for the variety. This is why prices in consumer price indices include VAT. In terms of such indices, P_{ijt} is referred to as the "purchaser price" (Eurostat, 2009).

The aggregate price and demand levels are P_t and Y_t , respectively. The resulting demand curve for variety i and good j are the ones implied by the nested CES preferences and equal:

$$Y_{ijt} = \left(\frac{P_{ijt}}{P_{jt}} \right)^{-\eta_j} Y_{jt} \text{ and } Y_{jt} = \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} Y_t, \quad (3)$$

where:

$$P_t = \left[\sum_{j=1}^J P_{jt}^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}} \text{ and } P_{jt} = \left[\int_0^1 P_{ijt}^{1-\eta_j} di \right]^{\frac{1}{1-\eta_j}}. \quad (4)$$

Here, $\eta_j > 1$ is the elasticity of substitution between varieties for good j and ε is the elasticity of substitution between goods.¹⁷ The latter is the parameter we aim to estimate. Since the firm is of negligible size in terms of the supply of varieties of good j , its choice of P_{ijt} does not affect the price of good j (Dixit and Stiglitz, 1983).

The existence of a Value Added Tax means that the firm does not receive all the revenue generated at the price P_{ijt} . Instead, the VAT involves charging a tax, τ_j , as a fraction of the pre-tax price.¹⁸ In terms of our notation, this means that, after the payment of the VAT, the firm receives $P_{ijt}/(1 + \tau_j)$ in net revenue per unit sold. Consequently, the firm's per-period flow profits are given by:

$$\frac{P_{ijt}}{1 + \tau_j} Y_{ijt} - W_t L_{ijt}. \quad (5)$$

To show that our results do not depend on whether or not one assumes price stickiness, we solve the firm's price setting decision under sticky prices, in a similar way to Calvo (1983). We then show that the price-stickiness parameters do not affect the relevant reduced-form elasticity of a good's price with respect to its VAT rate.

We assume that, in each period, with probability ϕ_j the firm can adjust its price costlessly, while with probability $(1 - \phi_j)$, it faces an infinite adjustment cost and keeps its price P_{ijt} fixed. The flexible price case is simply nested in this model and corresponds to the case in which $\phi_j = 1$.

The solution to this problem is such that the fraction ϕ_j of firms that reset their price will all choose the same price. This reset price, P_{jt}^* is proportional to a weighted average of the future marginal costs the firms face over the horizon that they have not adjusted their price and are still charging the reset price they currently choose. The proportionality factor is made

¹⁷Throughout this derivation we abstract from different expenditure weights across goods j . We do so to simplify notation. The main reduced-form equation we derive does not depend on this assumption.

¹⁸Because we focus on changes in steady-state relative prices, throughout, we assume that the VAT rates are constant at their steady-state values. Results with time-varying VAT rates are algebraically more cumbersome but yield the same elasticity as we derive here.

up of two components. The first is the gross markup factor $\left(\frac{\eta_j}{\eta_j-1}\right)$. The second is the gross VAT rate, $(1 + \tau_j)$. That is,

$$P_{jt}^* = \left[(1 + \tau_j) \left(\frac{\eta_j}{\eta_j - 1} \right) \sum_{s=0}^{\infty} \omega_{jt,s} MC_{jt+s}^* \right], \quad (6)$$

where MC_{jt+s}^* is the marginal cost at time $t + s$, given in equation (2), evaluated at the reset price P_{jt}^* . The weight, $\omega_{jt,s}$, is given by

$$\omega_{jt,s} = \left[(1 - \phi_j)^s \left(\prod_{j=0}^s \frac{1}{1 + r_{t+j}} \right) P_{jt+s}^{\eta_j} Y_{jt+s} \right] / \left[\sum_{q=0}^{\infty} (1 - \phi_j)^q \left(\prod_{j=0}^q \frac{1}{1 + r_{t+j}} \right) P_{jt+q}^{\eta_j} Y_{jt+q} \right]. \quad (7)$$

Because MC_{jt+s}^* is itself a function of the reset price, P_{jt}^* , equation (6) needs to be solved for P_{jt}^* to obtain the reset price. Doing so yields

$$P_{jt}^* = \left[(1 + \tau_j) \left(\frac{\eta_j}{\eta_j - 1} \right) \left(\frac{1}{1 - \alpha} \right) \sum_{s=0}^{\infty} \omega_{jt,s} W_{t+s} P_{jt+s}^{\frac{\alpha}{1-\alpha} \eta_j} \left(\frac{Y_{jt+s}^{\frac{\alpha}{1-\alpha}}}{A_{t+s}^{\frac{1}{1-\alpha}}} \right) \right]^{\frac{1}{1 + \frac{\alpha}{1-\alpha} \eta_j}}. \quad (8)$$

Based on this result, it is tempting to conclude that because for all producers of varieties of good j it is the case that

$$\frac{\partial \ln P_{ijt}^*}{\partial \ln (1 + \tau_j)} = \frac{1}{1 + \frac{\alpha}{1-\alpha} \eta_j}, \quad (9)$$

the elasticity of the price of good j , P_{jt} , after all firms adjust their price with respect to the value added tax rate is equal to the right-hand side of the above equation. However, this ignores the fact that, everything else equal, consumers will substitute away from goods whose value added taxes increase more than others. Therefore, in order to fully understand the effect of the VAT change on the price of a good, we have to solve for this substitution in demand. We do so below, under the assumption that the economy is in steady state, or rather, on a balanced growth path.

2.2 Relative prices in steady state

The balanced growth path is characterized by the following four properties: (i) Aggregate output Y_t grows at a constant rate, g , which is equal to the steady-state level of productivity growth $\frac{A_{t+1}}{A_t} = (1 + g)$. (ii) Inflation is constant, such that the aggregate price level, P_t , as

well as the prices of each of the goods, P_{jt} , grow at rate π .¹⁹ (iii) The real interest rate, r_t , is constant and equal to r . (iv) Nominal wage growth is constant and equal to productivity growth plus inflation, *i.e.*, $\frac{W_{t+1}}{W_t} = (1 + g)(1 + \pi)$.

On this balanced growth path, the forward-looking components of the reset price, P_{jt}^* , defined in equation (8), can be solved to obtain:

$$P_{jt}^* = \left[(1 + \tau_j) s_j W_t P_{jt}^{\frac{\alpha}{1-\alpha} \eta_j} \left(\frac{Y_{jt}^{\frac{\alpha}{1-\alpha}}}{A_t^{\frac{1}{1-\alpha}}} \right) \right]^{\frac{1}{1 + \frac{\alpha}{1-\alpha} \eta_j}}, \quad (10)$$

where the goods-specific constant, s_j , equals:²⁰

$$s_j = \frac{\eta_j}{(\eta_j - 1)} \frac{1}{1 - \alpha} \left(\frac{1 - \left(\frac{1 - \phi_j}{1 + r} \right) (1 + \pi)^{\eta_j} (1 + g)}{1 - \left(\frac{1 - \phi_j}{1 + r} \right) (1 + \pi)^{1 + \frac{\eta_j}{1 - \alpha}} (1 + g)} \right). \quad (11)$$

Given the Calvo-type price setting, the law of motion of the price level of good j , P_{jt} , as a function of the reset price, P_{jt}^* , and the previous period's price, P_{jt} , is

$$P_{jt} = \left[(1 - \phi_j) (P_{jt-1})^{1 - \eta_j} + \phi_j (P_{jt}^*)^{1 - \eta_j} \right]^{\frac{1}{1 - \eta_j}}. \quad (12)$$

On the balanced growth path, good j 's inflation rate equals π . In combination with the law of motion of the price level above, this allows us to solve for the level of the reset price set by the producers of varieties of good j that change their price, P_{jt}^* , relative to the overall price level, P_{jt} . This yields:

$$P_{jt} = P_{jt}^* F_j, \quad (13)$$

where:

$$F_j = \left[\frac{1 - (1 - \phi_j) (1 + \pi)^{\eta_j - 1}}{\phi_j} \right]^{\frac{1}{1 - \eta_j}}. \quad (14)$$

Combining this with the solution of the reset price from equation (10), we obtain that the

¹⁹If there are trends in relative prices, then there is neither a balanced growth path nor a steady state when $\varepsilon \neq 1$. The lack of a steady state and balanced growth path in this case is the main topic of studies of long-run structural transformation (see, for example, Herrendorf, Rogerson and Valentinyi, 2014). Because we are interested in shorter horizons, we abstract from such trends in relative prices in our derivations, which assures the existence of a balanced growth path. We do allow for such trends in our empirical analysis, however.

²⁰Our solution is derived under the assumption that $\left(\frac{1 - \phi_j}{1 + r} \right) (1 + \pi)^{1 + \frac{\eta_j}{1 - \alpha}} (1 + g) < 1$, which is true for common calibrations of r , π , g , η_j , and the price-stickiness parameter, ϕ_j .

relative price of good j along the balanced growth path is given by:

$$\frac{P_{jt}}{P_t} = (1 + \tau_j) F_j^{1+\frac{\alpha}{1-\alpha}\eta_j} s_j \frac{W_t}{P_t} \left(\frac{Y_{jt}^{\frac{\alpha}{1-\alpha}}}{A_t^{\frac{1}{1-\alpha}}} \right). \quad (15)$$

The final step in our derivation of the main equation for the steady-state relative price level is to substitute in the demand function for good j , equation (3), to take into account that shifts in the relative price, $\frac{P_{jt}}{P_t}$, affect the level of demand for good j , Y_{jt} . Doing so, yields that, in steady state, the relative price of good j equals:

$$\frac{P_{jt}}{P_t} = (1 + \tau_j)^{\frac{1}{1+\frac{\alpha}{1-\alpha}\varepsilon}} F_j^{\frac{1+\frac{\alpha}{1-\alpha}\eta_j}{1+\frac{\alpha}{1-\alpha}\varepsilon}} s_j^{\frac{1}{1+\frac{\alpha}{1-\alpha}\varepsilon}} \left[\frac{W_t}{P_t} \left(\frac{Y_t^{\frac{\alpha}{1-\alpha}}}{A_t^{\frac{1}{1-\alpha}}} \right) \right]^{\frac{1}{1+\frac{\alpha}{1-\alpha}\varepsilon}}. \quad (16)$$

Therefore, conditional on the real wage, $\frac{W_t}{P_t}$, the aggregate productivity level, A_t , and the level of output, Y_t , the elasticity of the relative price level of good j with respect to the VAT rate, τ_j , which we denote by β , is equal to:

$$\beta = \frac{1}{1 + \frac{\alpha}{1-\alpha}\varepsilon}. \quad (17)$$

Of course, this does not take into account that a VAT change potentially also has an effect on the overall economy, and thus on the real wage and output. However, the effect of the changes in these aggregate variables are the same across all goods j , which is what we exploit in the construction of the reduced-form equation that is at the heart of our empirical analysis.

The relationship between β and ε implied by equation (17) hinges on the assumption that firms face decreasing returns to scale production technology, *i.e.* $\alpha \in (0, 1)$. If $\alpha = 0$ then the production function, given in equation (1), has constant returns to scale and marginal costs do not vary with the level of output of the firm. Consequently, in this case, equation (6) implies that changes in VAT rates are fully passed through in prices and $\beta = 1$. If firms face decreasing returns to scale, their level of marginal costs depends on output. Hence, a change in the relative VAT rate of a good results in a relative price change that affects relative demand, and because of the decreasing returns to scale with respect to the flexible factor, it results in a change in the marginal cost of production. The equilibrium outcome is the fixed point in which the relative price change is consistent with the change in the marginal cost. This results in $\beta < 1$, which turns out to be what we find in the data.²¹

²¹Under our model assumptions, $\alpha = 0$ implies $\beta = 1$. This holds because we also assume that labor is mobile across all firms and sectors. Alternatively, if labor inputs were sector-specific, marginal costs would vary with

In addition, the CES preferences and [Dixit and Stiglitz \(1983\)](#) assumption about monopolistic competition imply that our model is derived under the assumption of constant gross markups over marginal costs (equation (6)). Long-run movements in gross markups in response to changes in relative VAT rates, for example because of entry and exit as in [Jaimovich \(2007\)](#), would affect the elasticity β . If, for example, gross markups permanently decline in response to an increase in the relative VAT rate of a good, then this would bias our estimate of β downwards and our estimate of ε upwards.

2.3 Reduced-form equation

To construct the reduced-form equation implied by the expression for the steady-state levels of the relative prices as in equation (16), we define the log of the price level of good j and the average log price across all goods as

$$p_{jt} = \ln P_{jt} \text{ and } \bar{p}_t = \frac{1}{n} \sum_{j=1}^n p_{jt}, \quad (18)$$

respectively.²² Using this notation, and the approximation $\ln(1 + \tau_{jt}) \approx \tau_{jt}$, we can rewrite equation (16) to obtain:

$$p_{jt} \approx \beta (\tau_{jt} - \bar{\tau}_t) + \delta_j + \xi_t + \bar{p}_t, \quad (19)$$

where:

$$\bar{\tau}_t = \frac{1}{n} \sum_{j=1}^n \tau_{jt}, \quad (20)$$

$$\delta_j = \beta \left[\left(1 + \frac{\alpha}{1-\alpha} \eta_j \right) \left(\ln F_j - \frac{1}{n} \sum_{k=1}^n \ln F_k \right) + \left(\ln s_j - \frac{1}{n} \sum_{k=1}^n \ln s_k \right) \right], \quad (21)$$

and

$$\xi_t = \frac{1}{1 + \frac{\alpha}{1-\alpha} \varepsilon} \ln \left(\frac{W_t}{P_t} \left(\frac{Y_t^{\frac{\alpha}{1-\alpha}}}{A_t^{\frac{1}{1-\alpha}}} \right) \right). \quad (22)$$

In practice, however, we do not have data for the log of the price *levels*, p_{jt} . So, empirically implementing the above as a reduced-form equation is not feasible. Instead, we have data on log price *indices*, changes in which are constructed to be proportional to changes in the log of the price levels.

Hence, to operationalize equation (19) as a reduced-form relationship, we focus on the

the level of output even under constant returns to scale. The latter case would also be consistent with $\beta < 1$.

²²Note that the log of the aggregate price level, P_t is p_t , not \bar{p}_t .

change in the log of the steady-state price levels, Δp_{jt} in response to a change in the VAT rate, $\Delta \tau_{jt}$ compared to the change in the average VAT across goods, $\Delta \bar{\tau}_t$. That is, the reduced-form equation that forms the basis of our empirical analysis equals:

$$\Delta p_{jt} \approx \beta (\Delta \tau_{jt} - \Delta \bar{\tau}_t) + \Delta \xi_t + \Delta \bar{p}_t. \quad (23)$$

Note that in this specific equation, the reduced-form parameter, β , represents the response of relative prices to changes in relative VAT rates, which from equation (17), only depends on two parameters: (i) the output elasticity of labor, $(1 - \alpha)$, and (ii) the between-goods elasticity of substitution, ε . The latter is the parameter we aim to estimate.

More important is the list of other parameters that it does *not* depend on. First, because we focus on steady-state levels of relative prices, β does not depend on the frequency of price adjustment. It is the same, no matter whether prices are flexible (*i.e.*, $\phi_j \rightarrow 1$) or sticky.²³ Second, neither heterogeneity in the degree of price stickiness, by assuming ϕ_j is different across goods (as in [Carvalho, 2006](#)), nor in markups, by assuming η_j varies across goods (as in [Eusepi, Hobijn and Tambalotti, 2011](#)), affect the response of relative prices to VAT rate changes.

One potential source of heterogeneity across goods that would affect β is heterogeneity in the output elasticity of labor, $(1 - \alpha)$. Such heterogeneity would result in different labor shares in the production of different consumption goods. However, [Fisher \(1969\)](#) shows that such heterogeneity would prevent us from finding a simple closed-form solution. Therefore, we abstract from this source of cross-good heterogeneity.²⁴ To put this potential heterogeneity in context, note that our production function is defined as an aggregate production function for a particular final good or service category bought by households for consumption purposes, and therefore, it does not include intermediate inputs.²⁵

2.4 Why not a general equilibrium analysis?

As discussed above, the reduced-form parameter β in equation (23), is the same under a broad set of underlying assumptions. These assumptions cover a broad class of models commonly used in macroeconomics.

²³We have derived our results under [Calvo](#)-style nominal rigidities. Our steady-state results are also valid under state-dependent pricing. See [Klenow and Kryvtsov \(2008\)](#) for a detailed comparison of models under these two types of price setting.

²⁴In addition, with identical growth rates of total factor productivity, g , a balanced growth path does not exist when the output elasticity of labor, $(1 - \alpha)$, varies across goods and $\varepsilon \neq 1$.

²⁵The parameter $(1 - \alpha)$ is the labor share over the whole (domestic) supply chain of these goods and services and the use of intermediate inputs is aggregated out using input-output relationships.

However, every estimation method is based on some identifying assumptions. The three that we make here are: (i) labor shares are the same across the production of consumer goods and services, (ii) labor is mobile across sectors, and (iii) markups are constant and determined by the CES preferences.

Though these identifying assumptions have their limitations, they are less restrictive than those that are made in other empirical strategies. The most common alternative method would be to estimate the parameters of the model, including ε , using a dynamic stochastic general equilibrium model. This is what [Karadi and Reiff \(2007\)](#) do using data for Hungary, which allows for the estimation of all the parameters, and not only ε , underlying the general equilibrium structure of the model.

Though such an approach allows one to focus on a broad set of parameters, it does require one to make specific assumptions about the sources of heterogeneity that our reduced-form parameter, β , does not depend on. Moreover, to close the model one also has to make specific assumptions about household preferences. In particular, about the intertemporal elasticity of substitution and Frisch elasticity of the labor supply. In order to fit the path of aggregate inflation, one also has to add a monetary policy rule.

Our approach, which identifies the between-goods elasticity of substitution, ε , from the correlation between long-run changes in relative prices and changes in relative VAT rates, is valid for any type of aggregate household preferences and monetary policy rule.

3 Empirical implementation

Exploiting the insight about the relationship between the upper-level elasticity of substitution and the long-run response of relative prices to changes in relative VAT rates in practice requires mapping equation (23) into existing data. In this section we describe how equation (23) can be estimated using a relatively simple panel-data regression that implements a local projection method ([Jordà, 2005](#)). This allows us to estimate the long-run response of relative prices to VAT changes.

3.1 Data

In principle, the elasticity, β , in equation (23) could be estimated solely based on cross-good variation in changes in relative prices. However, because changes in VAT rates in a given country are relatively infrequent, it is useful to pool the regression across countries.

Thus, our panel data analysis uses three sources of variation for the estimation of β ; goods,

j , countries, c , and time, t . The log of prices, p_{jct} , are measured using the logarithm of the monthly Harmonised Index of Consumer Prices (HICP).²⁶ We use them to calculate inflation at the good-specific level for 25 European Union (E.U.) countries in our sample. Our data cover the time period from January 1996 through January 2015.²⁷

Expenditures included in HICPs are classified in categories/goods, j , called COICOPs.²⁸ The COICOP classification system consists of three levels of aggregation. In our sample, the highest level of aggregation (one-digit level) consists of 12 *divisions*, such as food and non-alcoholic beverages, communication, restaurants and hotels, etc.²⁹ The next level of aggregation (two-digit level) is called a *group*. As an example, accommodation services is a group within the division of restaurants and hotels. Our sample includes 36 groups. The lowest level of aggregation (three-digit level) is a *class*. The group of alcoholic beverages, for example, consists of classes that cover spirits, wine, and beer, separately. Our sample includes 74 classes.

The level of granularity at which goods and services are defined matters for their substitutability. Hence, we report our estimate of β for each of these three different levels of aggregation at which we have data.

Data on VAT rates by COICOP for the 25 countries in our sample are not readily available. We construct them from two administrative sources: [European Commission \(2015\)](#), and [Eurostat \(2015\)](#). These give us information about which VAT rates are applicable to which goods and services in each country over time.

VAT rates are not the same for all goods within a country. Most countries have four different VAT rates that apply to different goods: super-reduced, reduced, standard, and parking rate. In addition, many countries have goods and services, such as education, for example, that are exempt from value added taxes.

E.U. law requires that the standard VAT rate is at least 15% and the reduced rate at least

²⁶The HICP is the euro area equivalent of the U.S. CPI, except that it excludes owner-occupied housing and includes rural consumers. Both the HICP and the CPI are constructed to approximate the increase in households' cost of living and are consistent with the CES preferences used in this paper.

²⁷The Appendix Table [A1](#) provides a complete list of countries in our sample. In particular, our sample includes class-level HICP price indices and VAT rates for Austria, Belgium, Cyprus, Czech Republic, Denmark, Germany, Estonia, Greece, Spain, Finland, France, Hungary, Ireland, Italy, Lithuania, Latvia, Luxembourg, Malta, Netherlands, Poland, Portugal, Slovakia, Slovenia, Sweden, and the United Kingdom. Austria, Denmark and Sweden have no documented VAT rate changes between 1996 and 2015. In addition, our data originally included Bulgaria and Romania, which we dropped because those countries faced periods of hyperinflation.

²⁸COICOP is an acronym for "Classification Of Individual Consumption accOrding to Purpose."

²⁹The complete list of COICOPs is provided in the Appendix Tables [A2](#) to [A4](#). Our initial data include 12 divisions, from which we dropped all classes (and groups) pertaining divisions 6 (Health) and 10 (Education) because of the non-market nature of price setting in these sectors. In addition, our data do not include two additional divisions that cover spending by non-profit institutions serving households (NPISH) and government consumption for which price data is imputed rather than directly measured.

5%. Actual rates applied vary across countries.³⁰ For example, in 2015 goods facing a standard rate in the United Kingdom were charged a 20% VAT rate, while in Luxembourg the standard rate was 15%. In addition, the four broad rate categories can also include a range of levels of VAT rates, and the level applied to a certain category may change over time.³¹ Sometimes a good or service gets reclassified into a different VAT rate category. For example, before September 2012 cosmetic surgery in Spain was charged a reduced VAT rate while afterward it fell under the standard VAT rate category. Finally, the same categories may face different VAT rates in different countries. The result is that VAT rates for the same goods and services vary across countries and over time. For example, as of January 2015, Spain charged a 10% VAT rate on hotel accommodation services while Portugal charged a 6% rate. Restaurants faced a 7% VAT rate in France and a 19% rate in Germany.

The most important variation in VAT rates for the estimation of β , however, is changes in VAT rates on specific goods and services over time. Most of these changes are because the rate associated with the VAT category that a good or service is classified in changes. For example, between June 2010 and September 2012, Spain increased its standard VAT rate from 16% to 21%. The majority of these VAT rate changes in our sample occur either on January 1st or July 1st.

Our matching of the administrative data on VAT rates with COICOPs yields good-country-specific time series for the applicable VAT rate, τ_{jct} , where j is defined, alternatively, at the COICOP class, group, or division levels. VAT rates for the group- and division-level expenditure categories are constructed as weighted averages of the VAT rates in the underlying classes.³²

The HICP price data, p_{jct} , and VAT data, τ_{jct} , are, respectively, the left- and right-hand side variables of our reduced-form equation (23). What is left is to map equation (23) into a specification that allows for the identification and estimation of β for the country-COICOP-time panel structure of our data.

³⁰These simple rules are, however, complicated by a multitude of derogations granted to certain European Union Member States, and in some instances, to a majority of Member States.

³¹For example, in 2010, France applied two different super-reduced VAT rates, 2.1% and 5.5%. In 2012, a third 7% super-reduced rate was introduced.

³²Due to lack of data on consumer expenditures at the class level, to aggregate from the class up to the group level, we use an equally weighted average of VAT rates across classes within each group. We use consumer expenditure shares to aggregate from group up to division level. The construction of our data is described in more detail in Appendix A.1. In order to minimize the measurement error in the VAT rate changes for higher level of aggregation, we only include observations in our data if the within-category dispersion of VAT rates is smaller or equal to one percentage point. This is to avoid attenuation bias from affecting our results.

3.2 Model specification and identification

Equation (23) describes log changes in relative price levels, for good j , between two steady states that differ in the VAT rates charged. In practice, of course, log changes in observed prices reflect more than only shifts between steady states. Throughout, we interpret the steady-state response of relative prices as *long-run* movements in the data.

In particular, we estimate β as the effect of cumulative VAT rate changes from t to $t+l$, which we denote by $\Delta_l \tau_{jct+l} = \tau_{jct+l} - \tau_{jct}$, on the cumulative log change of prices from t to $t+h$, where $h \gg l$. We denote this cumulative log change in prices by $\Delta_h p_{jct+h} = \ln p_{jct+h} - \ln p_{jct}$.

For each choice of the length of period over which we accumulate VAT changes, l , and the horizon over which we consider the long-run response of log prices, h , we obtain an estimate $\beta_{l,h}$. This parameter is estimated using the local projection method introduced in Jordà (2005).

For the particular estimation problem at hand, this amounts to running a panel data regression of the form:

$$\Delta_h p_{jct+h} = \beta_{l,h} \Delta_l \tau_{jct+l} + \gamma_{ct} + \alpha_{jcm} + u_{jct}. \quad (24)$$

In addition to the changes in the log prices and VAT rates, this equation contains a country-time fixed effect, γ_{ct} , and a COICOP-country-month fixed effect, α_{jcm} . u_{jct} is the residual.

The country-time fixed effect absorbs both the effect of the average change in VAT rates, $\Delta \bar{\tau}$, the average log change in prices, $\Delta \bar{p}_t$, as well as the change in country-wide economic conditions, $\Delta \xi_t$, from equation (23). Including this fixed effect means that we do not have to specify our regression in terms of the logs of relative prices and relative VAT changes. The country-time fixed effect captures the country-specific changes in the log of the overall price level and the average VAT rate change across goods in a country. We also include a country-COICOP-month fixed effect, α_{jcm} , to allow for potential trends in relative prices across goods and countries, as well as seasonal effects in country-good-specific inflation rates.

Figure 1 helps illustrate our identification strategy. We estimate the effect of tax changes between t and $t+l$, *i.e.*, $(\tau_{t+l} - \tau_t)$, on the cumulative log price change between $t+h$ and t , $(p_{t+h} - p_t)$, where $h \gg l$.

For our estimate of $\beta_{l,h}$ to be consistent, the VAT changes, $\Delta_l \tau_{jct+l}$, need to have three properties. First, they need to be uncorrelated with the grayed out future tax changes, between $t+l$ and $t+h$, in Figure 1. If this is not the case, then our estimate of $\beta_{l,h}$ suffers from omitted variable bias, since it will partly include the effect of tax changes between $t+l$ and $t+h$, rather than between t and $t+l$, on the log change of the price level. Secondly, since we focus on the long-run response of prices to the VAT changes, the VAT changes, $\Delta_l \tau_{jct+l}$, need to be (expected to be) permanent. If this is not the case then the long-run response to the VAT

changes we identify will be muted in the data compared to our model. Finally, they need to be exogenous with respect to the residual u_{jct} . We revisit these three properties in Subsection 4.1, where we provide evidence in support of them.

In our derivation of equation (23) we use labor as the only adjustable factor of production. The parameter α reflects the degree of decreasing returns to scale of the production function with respect to this factor. Because labor is mobile across firms and sectors, all firms pay the same real wage and face the same marginal cost schedule. Our empirical approach in equation (24), however, allows for marginal cost schedules to vary across goods. As long as changes in these schedules are uncorrelated with VAT changes, $\Delta_l \tau_{jct+l}$, our estimate of $\beta_{l,h}$ remains consistent. In that case, these orthogonal changes in marginal costs generate unexplained changes in log prices, $\Delta_h p_{jct+h}$, that are either absorbed by the fixed effect, α_{jcm} , or are part of the residual, u_{jct} .

For the practical implementation of equation (24) it remains to choose l and the value of h that we interpret as the “long-run.” As for the value of l , we go with the natural choice of $l = 12$. That means we consider the long-run effect of 12-month changes in VAT rates, $\Delta_{12} \tau_{t+12}$, on log price levels.

What exactly we mean by “long-run” is less clear-cut. The obvious choice seems to associate long run with h as large as possible. In practice, however, increasing h reduces the effective sample size and thus the degrees of freedom and the precision of the parameter estimate we are interested in. Moreover, choosing h too large for the estimation of equation (24) poses a theoretical challenge to our approach.

This challenge is that the production function in equation (1) has decreasing returns to scale in the adjustable production factor, which is labor in this case. Therefore, implicitly, it assumes that capital inputs are fixed. The decreasing returns to scale with respect to the adjustable inputs are captured by the term $\frac{\alpha}{1-\alpha}$ in the expression for the reduced-form parameter β in equation (17). Therefore, in the longer run, the appropriate standard production function would also include capital as an adjustable factor, making returns to scale constant. However, solving the model under constant returns to scale to adjustable inputs results in $\beta = 1$, which does not depend on ε . As we show later, however, $\beta < 1$ in the data.

Though it might seem contradictory to estimate “long-run” effects assuming the capital inputs are fixed, the relevant “long-run” for our purpose is the duration of the transitional price setting dynamics in response to a change in VAT rates. Such transitional dynamics tend to die out after about 40 months in common multisector New Keynesian models.³³ In the short

³³See [Carvalho \(2006\)](#), [Bouakez, Cardia and Ruge-Murcia \(2009\)](#), and [Carvalho and Nechio \(2016\)](#) for a detailed analysis of such transitional dynamics.

run, the transitional price-setting dynamics in these models are affected by the degree of price stickiness across goods and services, which is something that the steady-state effects that we aim to estimate does not depend on. Thus, choosing h too small would render our estimates inconsistent. With this in mind, we focus on $h = 48$ in our baseline set of results and discuss how they change when we vary h between 12 to 62.

4 Empirical results

We present our empirical results in four parts. First, we document the amount of variation in VAT rates across countries and COICOPs. This is the variation in $\Delta\tau_{jct}$, on the right-hand side of equation (24), that we use to identify the parameter $\beta_{l,h}$. Next, we present the results for our baseline specification for the lowest level of aggregation of COICOPs, where goods and services are defined at the class level. Third, we show that these results are robust to various different model specifications and sample choices. Finally, we present the results for the other (higher) levels of aggregation, namely at the group and division levels.

4.1 Changes in VAT rates

Of course, as with any regression, our estimation, based on equation (24), needs substantial variation in the explanatory variable of interest to reliably estimate the associated coefficient. Hence, before we present our estimation results, we consider the variation in changes in VAT rates in our sample. We provide evidence that these changes (*i*) are uncorrelated with future changes in VAT rates, (*ii*) are expected to be permanent, and (*iii*) that the variation in these changes is exogenous with respect a wide range of factors affecting long-run price changes.

Variation in VAT changes

Figure 2 provides two measures of the variation in VAT rate changes across countries and COICOPs at the class level of aggregation. The first is the monthly count of the number of non-zero 12-month changes in VAT rates across COICOPs and countries. This is depicted by the dark-shaded area in the figure. Changes in VAT rates have occurred during the whole sample period, although they are concentrated in the post-2008 part of the sample. This means that the bulk of the variation that identifies our parameter of interest, $\beta_{l,h}$, is in the last seven years of the sample.

Note, however, that while the right-hand-side variable of equation (24) corresponds to a simple 12-month change in the VAT rate by country and COICOP, our estimates reflect the

effects of VAT rate changes on prices after accounting for COICOP-country-month (α_{jcm}) and country-time (γ_{ct}) fixed effects. Therefore, the variation that we actually exploit in our empirical approach involves *demeaned* changes in VAT rates that result from the absorption of these fixed effects. This demeaned variable captures deviations of changes in VAT rates from the average change over 12-months across COICOPs in a country.

The second measure of variation in VAT rate changes in Figure 2, depicted by the light-shaded area, is the monthly standard deviation of this demeaned variable across countries and COICOPs. This measure shows that, while the concentration of changes in VAT rates is in the latter part of the sample, the effective variation in the explanatory variable, after taking into account the fixed effects, is more evenly distributed across the sample.

Hence, our sample to estimate $\beta_{l,h}$ includes not only a large number of VAT changes, but also substantial variation in changes in VAT rates across COICOPs, countries and time.³⁴

Exogeneity and expected permanence of relative changes in VAT rates

When we introduced our identification strategy in Section 3.2, we emphasized that, for us to obtain a consistent estimate of β , the VAT changes ($\Delta_l \tau_{jct+l}$) need to be (i) uncorrelated $\Delta_{h-l} \tau_{jct+h}$, (ii) (expected to be) permanent, and (iii) exogenous with respect to the residual u_{jct} .

To test whether future changes in tax rates are uncorrelated with the changes we include as explanatory variables, we estimate the effects of $(\tau_{t+h} - \tau_{t+l})$ on $(\tau_{t+l} - \tau_t)$ by performing simple OLS regressions for $l = 12$ and $h = 12, \dots, 62$. We find that changes in tax rates between $t + l$ and $t + h$ are only very weakly correlated with changes in taxes between t and $t + l$, with all regressions yielding a $R^2 < 0.002$.³⁵

Moreover, in line with the second property, the changes in VAT rates included in our data set are persistent. A simple Ordinary Least Squares (OLS) regression of VAT rates on their lags, controlling for the same fixed effects included in equation (24), shows that the coefficient on the first lag is near one (specifically, 0.997), and it is the only statistically significant coefficient.³⁶

As for the exogeneity of the relative VAT changes in equation (24), it is important to realize that most factors with which VAT rate changes are potentially correlated are captured by the fixed effects. Inclusion of the country-time fixed effects, γ_{ct} , means that our estimates are not affected by any country-specific factors that vary over time, like country-specific business

³⁴More detailed summary statistics about the incidence of and variation in VAT changes across countries and COICOPs can be found in Appendix A.1.

³⁵This also holds when we consider other levels of aggregation.

³⁶In particular, we consider the regression: $\tau_{jct} = \sum_{k=1}^{24} \lambda_k \tau_{jct-k} + \gamma_{ct} + \alpha_{jcm} + \varepsilon_{jct}$. Estimates of this equation by country yield qualitatively similar results.

cycles, legislation, or inflationary effects. In addition, the COICOP-country-calendar-month fixed effect, α_{jem} , filters out seasonal fluctuations by COICOP and country. Therefore, this leaves only the case in which the incidence of legislated relative VAT changes across goods is correlated with factors that affect long-run relative price changes across goods beyond the VAT rate changes themselves.

In the robustness checks of our results we consider alternative specifications of the fixed effects that soak up country-group- and country-division-wide omitted variables that affect both relative VAT changes and relative prices. Such variables can be the source of these correlations that would violate our exogeneity assumption. As discussed in that section we do not find that such variables have a substantial impact on our results.

4.2 Results at class level

The dark line in Figure 3 depicts the estimates $\hat{\beta}_{l,h}$ for $l = 12$ and $h = 12, \dots, 62$. The shaded area is the associated 95 percent confidence interval. The estimates of β initially decline as a function of h , from 0.30 at $h = 12$ to 0.24 at $h = 25$. After that, $\hat{\beta}_{l,h}$ steadily increases and peaks at 0.43 at $h = 52$.

At our choice for the long run, $h = 48$, the estimated effect of changes in VAT rates on long-term changes in relative prices equals 0.40. Column (I) of Table 1 provides the detailed regression results for this baseline specification. The estimate of β is relatively precise with a standard error of 0.04.

To map the estimated coefficient $\hat{\beta}_{l,h}$ into the implied elasticity of substitution, $\hat{\varepsilon}$, using equation (17), one needs to take a stand on the value of α . The value of α is commonly picked based on evidence on the average labor share, which equals $(1 - \alpha)$ in the class of models we consider. European data on GDP and labor compensation suggest that, despite some variability across countries, the average labor share for the European Union is approximately 2/3 (European Commission, 2007, Chapter 5). In line with this evidence, we calculate our baseline results for $(1 - \alpha) = 2/3$. In Subsection 4.3, we specifically consider the sensitivity of our results to varying α to a degree consistent with the cross-good labor share variations in the U.S., reported in Eusepi, Hobijn and Tambalotti (2011, Table 1). We focus on this estimated variation because it calculates the labor shares for the production of consumer goods and services and aggregates out intermediate goods and services purchases. This aggregation results in much less variation in labor shares across consumer goods and services (COICOP categories) than one observes across industries.

As can be seen from the bottom two rows of Column (I) of Table 1, the elasticity of

substitution, $\hat{\varepsilon}$, implied by $(1 - \alpha) = 2/3$ and our parameter estimate $\hat{\beta}_{l,h} = 0.40$, is 2.98. Taking into account the standard error of the reduced form parameter, this suggests that, with a 95% probability, the implied elasticity of substitution is between 2.1 and 4.3. Thus, our evidence at the class level of aggregation of COICOPs is that goods and services bought by consumers at this low level of aggregation are substantially more substitutable than in case of Cobb-Douglas preferences. However, the elasticity of substitution is much lower than 11, which is the value of the between-varieties elasticity of substitution, η_j , that most commonly gets used.³⁷

4.3 Robustness of class-level results

The value of the implied elasticity is robust with respect to many perturbations of our sample and model. To illustrate this, in this subsection we present different estimates of β for a shortened sample period, for a different specification that includes lagged inflation, for different choices of h and l , and for different subgroups of goods and services. In addition, we also discuss how the choice of α affects our results, why it is important to pool the model across many countries and COICOPs, why we conclude there are little or no anticipatory price increases in our data. Finally, we check for the robustness of standard errors and filter out potential omitted variables.

Sample period and lagged inflation

Table 1 reports estimates for an alternative sample period in Column (II), and estimates for a specification that includes lagged inflation in Column (III). The results reveal that the estimated reduced-form parameter, $\hat{\beta}_{12,48}$, as well as the implied elasticity of substitution, does not vary much when we change the sample period or add lagged inflation terms to the model.

In particular, Column (II) of Table 1 lists the results obtained by estimating equation (24) for the time period that *excludes* the Great Recession and its aftermath. The point estimate of the reduced form parameter, β , is almost the same if one excludes the VAT changes in the latter seven years of the sample. However, because most of the variation in VAT rates occurs in these years (see Figure 2), the exclusion of these dates widens the confidence interval for the implied elasticity.

Column (III) of Table 1 contains the results for an extension of equation (24) that includes 12 lags of COICOP-country-specific inflation. We include these lags of inflation to capture

³⁷This common calibration of η_j is based on evidence in Basu and Fernald (1997) that suggests that markups of price over marginal cost are in the order of 10 percent.

potential inflation dynamics at work at the time of the VAT rate changes that might influence our estimated β because of omitted variable bias. Thus, the generalized specification estimated reads

$$\Delta_h p_{jct+h} = \beta_{l,h} \Delta_l \tau_{jct+l} + \sum_{k=1}^{12} \psi_k \Delta p_{jct-k} + \gamma_{ct} + \alpha_{jcm} + u_{jct}, \quad (25)$$

where ψ_k for $k = 1, \dots, 12$ are the added parameters.

It turns out that adding these lagged inflation terms does not substantially alter our estimates. The implied elasticity in this generalized specification is 2.3 with an associated 95% confidence interval ranging from 1.7 to 3.1. Again, this suggests that at this low level of aggregation consumer goods and services are more substitutable than Cobb-Douglas but a lot less than the commonly assumed substitutability between varieties.

Sensitivity to choice of h and l

We already discussed how theoretical results guided us to choose $h = 48$ as our benchmark and showed in Figure 3 that, over the range $h = 36, \dots, 60$, our estimates for $\hat{\beta}_{12,h}$ are very similar.

The blue lines in the same figure provide more evidence in support of our choice of $h = 48$. The estimation of equation (24) focuses on the effects of 12-month changes of in VAT rates. Alternatively, one could slice this 12-month change up into 12 one-month changes. This results in an unrestricted version of regression (24), in which, for each horizon h , we quantify the effect of all one-period changes in VAT between t and $t + 12$. The resulting regression specification is

$$\Delta_h p_{jct+h} = \sum_{m=1}^{12} \beta_{12,h,m} \Delta \tau_{jct+m} + \gamma_{ct} + \alpha_{jcm} + u_{jct}, \quad (26)$$

where $\Delta \tau_{jct+m} = \tau_{jct+m} - \tau_{jct+m-1}$.

Our theory suggests that β is the same no matter what the choice of the parameter l over which we measure VAT rate changes, as long as we focus on their long-run effects. So, splitting up our baseline choice of $l = 12$ into its 12 subperiods should yield that $\beta_{l,h,m} = \beta_{l,h}$ for all $k = 1, \dots, 12$.

The blue lines in Figure 3 show the estimates $\hat{\beta}_{12,h,m}$ where $m = 1, \dots, 12$. As can be seen from the figure, there is noticeably more variation across m in the $\hat{\beta}_{12,h,m}$'s for $h < 24$ than for $24 \leq h \leq 50$. In fact, a formal F-test of $H_0 : \beta_{l,h,m} = \beta_{l,h}$ for all $k = 1, \dots, 12$ rejects this null at a 10 percent significance level for all $h < 22$. For all $h \geq 22$ this null is not rejected. In addition, the variation in $\hat{\beta}_{12,h,m}$'s does increase a lot after $h = 50$. At our choice of $h = 48$ the variation in the point estimates $\hat{\beta}_{12,h,m}$ across m is the smallest in our sample.

This result provides corroborating evidence in support of our choice of $h = 48$. Furthermore, because the estimated $\hat{\beta}_{12,h,m}$'s are so similar for the different m , the result also shows that our baseline result in Column (I) of Table 1 is not driven by our choice of $l = 12$.

Sensitivity to the choice of α

Our mapping between the long-term estimate of $\beta_{12,48}$ and the upper-level elasticity of substitution is directly dependent on our assumed value of the labor share. While our choice for the benchmark level of the parameter α was guided by the empirical evidence on the average labor share across goods in the U.S. and across E.U. countries, given the estimate of β reported in Table 1, a different choice of α implies a different value of ε .

To consider how the implied elasticity is affected by α , the blue line in Figure 4 reports how the estimated ε varies as the labor share $(1 - \alpha)$ increases from 0.4 to 0.85. The shaded area is the associated 95 percent confidence interval for the class-level estimates. We discuss the results for the other levels of aggregation in Section 4.4.

Figure 4 shows that, as the labor share increases, the point estimates for the elasticity ε increases. As the labor share increases from 66%, our benchmark level, to 85%, for example, point estimates for the upper-level elasticity of substitution increase from 3 to 8 at the class level. The larger the labor share, however, the smaller the precision of our estimates, as can be seen in the widening of the 95% confidence interval for the class level.

Because our baseline results are conditional on one particular value of α , we consider the sensitivity of our results to the uncertainty about α . For the class of goods in our data, U.S. evidence suggests the labor share varies between 52% and 77% (see Eusepi, Hobijn and Tambalotti, 2011, Table 1). If we take this evidence into account and assume that $(1 - \alpha) \sim Unif(0.52, 0.77)$, we obtain a point estimate of 2.91, instead of our baseline estimate of 2.98. However, the uncertainty about α widens the confidence interval to [1.4, 5.6].

One other interesting case to consider is the one where $\alpha = 0$ and the production function exhibits constant returns to scale. In that case equation (17) implies that $\beta = 1$ and there is full pass-through of VAT changes into prices. This, however, contradicts what we observe in the data where we obtain that $\beta < 1$.

The virtue of pooling across countries

Thus far, we focused on results obtained from the estimation of equation (24) with data pooled across both countries and COICOPs. Of course, the expression (23), on which equation (24) is based, is only for one particular country.

The reason we pool our results across countries is that VAT rate changes are infrequent enough that there is not enough variation in the right-hand-side variable to reliably estimate the reduced-form parameter, β , separately for all countries.

To illustrate this, we estimate the following country-specific regressions:

$$\Delta_h p_{jt+h}^{(c)} = \beta_{l,h}^{(c)} \Delta_l \tau_{jt+l}^{(c)} + \gamma_t^{(c)} + \alpha_{jm}^{(c)} + u_{jt}^{(c)}, \quad \forall c, \quad (27)$$

where γ_t is a time fixed effect and α_{jm} is a COICOP-month fixed effect.

Figure 5 shows the distribution of the country-specific estimators of $\hat{\beta}_{12,48}^{(c)}$, as well as the asymptotic distribution of the pooled estimate, $\hat{\beta}_{12,48}$, reported in Table 1.³⁸

The comparison between the two distributions highlights the gains from pooling across countries. First of all, many of the country-specific point estimates are negative which implies an implausible elasticity of substitution, ε , smaller than -0.5 . Moreover the standard errors of the country-specific coefficients are rather large.

The pooled OLS estimate, $\hat{\beta}_{12,48}$, is a complicated weighted average of the country-specific ones where those country-specific coefficients with the smallest estimated standard errors get the highest weights.³⁹ The pooled parameter estimate coincides with the only mode of the density of the country-specific $\hat{\beta}_{12,48}^{(c)}$'s in positive territory. Moreover, because it is based on many more VAT changes, it is estimated with a much smaller standard error than the country-specific parameters.⁴⁰

Possible anticipatory effects

Our estimate $\hat{\beta}_{l,h}$ implies a less than one-to-one pass-through from changes in VAT rates to prices. While this finding is not at odds with some of the literature that explored the effects of changes in VAT on inflation in specific countries, it is possible that our estimates do not capture anticipatory effects of changes in VAT rates.⁴¹ Such anticipatory effects occur when relative prices move in response to the announcement of VAT changes rather than in response to the actual VAT change that we measure in our data.

³⁸The distribution of the $\hat{\beta}_{12,48}^{(c)}$'s depicted in the figure is a kernel density estimate using a normal kernel and setting the bandwidth for each $\hat{\beta}_{12,48}^{(c)}$ equal to its estimated standard error.

³⁹See [Wooldridge \(2001\)](#), page 150, for a derivation of this result.

⁴⁰Though, a formal F-test for the null $H_0 : \hat{\beta}_{12,48}^{(c)} = \hat{\beta}_{12,48}$ for all c is rejected at a 1% significance level.

⁴¹In particular, [Andrade, Bénassy-Quéré and Carré \(2015\)](#) use firm-level French trade data to show that the degree of pass-through of changes in VAT rates to prices can be smaller than one-to-one, as it depends on both the degree of competition on a given market and each firm's market share. Using Hungarian micro-level data, [Karadi and Reiff \(2014\)](#) also find that the degree of pass-through can be smaller than one-to-one depending on the size and the sign of changes in VAT shocks. Using a dataset similar to ours, [Benedek, de Mooij and Wingender \(2015\)](#) also find incomplete pass-through of VAT changes to prices.

Changes in VAT rates are frequently announced in advance of their effective implementation. For example, on November 2008, the U.K. government announced a reduction in its standard VAT rate from 17.5% to 15% to become effective by December 2008 (see [Seely, 2013](#)). In addition, in a case study of a pre-announced tax hike in Germany, [Carare and Danninger \(2008\)](#) find some of the effects in inflation to occur during the announcement period. Therefore, one possible concern is that our estimates of $\beta_{l,h}$ do not account for the whole effect of VAT rate changes on prices.

Because we have no comprehensive data on announcement dates for all VAT rate changes included in our sample, we test whether prices react to changes in VAT rates before their effective implementation by estimating our main regression, equation (24), for h ranging between 1 and 12. This captures announcement effects in the following way. Suppose that VAT rates changed in July 2005. Then, our regression for $h = 1, \dots, 12$ considers the relative price change in August 2004 to be affected by the change in VAT rates over the subsequent 12 months that include July 2005. Hence, this regression captures anticipatory effects up to a year before the actual VAT change.

Our results provide little evidence of anticipatory effects of changes in VAT rates on prices. In particular, the estimated $\beta_{l,h}$ for $h = 1, \dots, 12$ roughly increases linearly up to its value at $h = 13$, before attaining the dynamics portrayed in [Figure 3](#). This pattern is consistent with no anticipatory effects at all. To see why this is the case, consider the following example.

Suppose VATs go up by one percent in January 2020 and prices do as well. In our regression, this change in VAT rates is captured in 12 monthly observations because we look at the 12-month change in VAT rates. So, the right-hand-side variable will be one for all months between January 2019 and January 2020.

What fraction of these observations match a change in the prices? This depends on the left-hand-side variable. At $h = 1$ the left-hand-side variable looks at the monthly price change, which will be zero in all observations except in January 2020. Thus, only one out of the 12 measured changes in the VAT rates corresponds to this price change. Consider two-month changes in prices. Then the left-hand-side variable will equal one percent in both January 2020 as well as December 2019. In this case, two observations on the price change match the positive 12-month changes in taxes. As a result, the estimated β will be twice as large. Hence, the estimated coefficient increases linearly between $h = 1$ and $h = 12$ if there are no anticipatory effects.

This is the reason that we present our results for $h > 12$ in our empirical analysis. For $h = 1, \dots, 12$ the estimated reduced-form coefficient approximately increases in such a linear fashion, which is consistent with the absence of significant anticipatory effects.

Possible upward bias on our estimate of ε

If there were anticipatory effects that we ignored in our analysis, then this would bias our estimate of β downwards and, because of equation (17), our estimate of ε upwards. Our estimate of ε is already relatively low, however. Such a bias would reduce the estimate even more.

There are two other potential sources of bias that would skew our results in the same direction. The first is that our estimates are derived under the assumption of constant gross markups, implied by the Dixit and Stiglitz (1983) setup we use. If gross markups actually absorb part of the VAT changes then this would reduce the relative-price response to change in the relative VAT rate. This would bias our estimate of the reduced-form parameter downwards. As we discussed in Subsection 2.2, in that case our estimate of ε is upward biased.

Similarly, if incorrect matching of VAT rates with COICOP categories or the aggregation of the VAT rates across categories introduced measurement error in $\Delta\tau_{jct}$ then our estimate of β would be downward biased as well. Again, this implies that our estimate of ε is upward biased.⁴²

Thus, if anything, our results indicate that the true upper-level of substitution is equal to or below our estimate of about 3.

Comparison with other elasticity estimates

To put our results in the context of existing estimates of substitution elasticities, it is important to realize that the upper-level elasticity of substitution that we estimate measures the willingness of consumers to substitute between different categories of final goods and services that they consume. For example, it represents households' willingness to substitute between haircuts and jewelry. We choose this example because haircuts are non-tradable services while jewelry consists of tradable goods.

Many studies report estimates of elasticities of substitution for tradable goods and services.⁴³ Though all of these parameters capture degrees of substitutability, they conceptually differ from our upper-level elasticity of substitution in the scope of the goods *between* which they capture substitutability. In particular, trade elasticities reflect how substitutable the same goods and services imported from different countries are, irrespective of whether these goods are bought by

⁴²In general, any attenuation bias that biases our estimate of β towards zero will bias our estimated elasticity, ε , upwards.

⁴³For example, Broda and Weinstein (2006), Simonovska and Waugh (2014a), Feenstra et al. (2014), Imbs and Mejean (2015), Simonovska and Waugh (2014b). Some of these studies report estimates of the Fréchet parameter based on Eaton and Kortum (2002). This parameter equals the elasticity of substitution minus one. Other studies directly report estimates of elasticities of substitution.

consumers, businesses, or the government. To put this in the context of our previous example, such a trade elasticity measures how quickly European consumers substitute between jewelry imported from India and jewelry imported from Brazil. Of course, studies of trade elasticities only focus on tradable goods and services and thus, do not include haircuts. However, even if we select a sample of tradable goods only, our estimated elasticity would pertain to the substitution between, for example, jewelry and books, and still conceptually differ from trade elasticities.

We are not the first to point out the conceptual difference between the upper-level elasticity of substitution that is the focus of our analysis and the trade elasticities estimated in many other studies. In fact, [Eaton and Kortum \(2002\)](#) explicitly distinguish between these two types of elasticities.⁴⁴

With the caveat of the conceptual difference between our elasticity of interest and the trade elasticity in mind, it is still useful to compare the magnitude of our estimated elasticity with those from the trade literature. To make the scope of goods and services in our analysis more comparable to those used in the trade literature, we separate items at the class level into tradables and nontradables using the classification from [Allington, Kattuman and Waldmann \(2005\)](#). Columns *(II)* and *(III)* in Table 2 list the results for these two categories. They show that the estimated reduced-form coefficients, as well as elasticities for tradables and nontradables, are very similar. Our point estimate of the elasticity of substitution between tradable consumer goods and services is 3.5 with a 95% confidence interval ranging from 2.2 to 5.9.

The many papers that estimate trade elasticities rely on different methodologies and vary in the goods and services as well as countries and periods covered. In an effort to report the estimates most comparable to our result, we focus on their estimates at levels of aggregation as close as possible to the 3-digit class level we use for the estimate we report in Column *(II)* of Table 2.

Probably the most comparable to our estimate is [Broda and Weinstein \(2006\)](#), who at the 3-digit level obtain a point estimate of the elasticity of 4. Several other studies have found aggregate trade elasticities of a similar order of magnitude. For example, [Imbs and Mejean \(2015\)](#) report an estimate of 3.9, [Simonovska and Waugh \(2014a\)](#) obtain 4.1, and [Caliendo and Parro \(2015\)](#) get 4.6.

Our estimated upper-level elasticity of substitution for tradable consumption goods and

⁴⁴In [Eaton and Kortum \(2002\)](#) the upper-level elasticity is defined as the parameter σ in their equation (3), while the trade elasticity is defined as the Fréchet parameter from the productivity distribution in their equation (4). [Broda and Weinstein \(2006\)](#) also discuss how the elasticity of substitution between tradable varieties that they estimate can be interpreted as a lower-level elasticity of substitution in a CES preference specification. That is, the elasticity of substitution they estimate is the equivalent of η_j rather than of ε in our analysis.

services is at the lower end of this range of estimates of elasticities of substitution from the trade literature. This is not surprising because these studies tend to focus on the degree of substitutability of more closely related varieties of goods and services than we do.

Clustered VAT changes and standard errors

The last concern about our identification strategy is the possibility that VAT changes happen in bunches, at the division or group level. In particular, these divisions and groups might have been singled out because of category-specific circumstances. This raises two potential issues for our empirical method.

The first is that such common circumstances mean that the standard errors that we present do not take into account that observations may be correlated at the division-country-date or group-country-date levels. To address this, Columns (IV) and (V) in Table 2 present our estimate with standard errors clustered at the group-country-date and division-country-date levels, respectively. They show that the clustered standard errors are only slightly bigger than those in the benchmark specification.

The second is that the circumstances that lead to the bunching of VAT changes at the division or group levels constitute potential omitted variables. In order to address this, we replaced the country-time dummy, γ_{ct} , in our baseline regression (24) with either a country-group-time dummy (γ_{cgt}), or a country-division-time dummy (γ_{cdt}). These dummies soak up the between-group and between-division variation in our data, respectively. Thus, the results from this alternative regression specification are solely due to the within-group and within-division variation in the data. As you can see from Columns (VI) and (VII) in Table 2, this generalized specification yields estimated reduced-form coefficients that are very similar to the ones reported for the baseline regression in the paper. From this we conclude that the parameter estimates we report for our baseline specification in the paper are mainly due to within-group and within-division variation in VAT changes in our data. Thus, the bunching of VAT changes does not seem to be a major driver of our results.

This is not completely surprising. The Appendix Table A2 suggests that there are many class-specific VAT changes in our dataset. And the regression results in Table 2 suggest that these provide a large part of the variation that identifies our parameter estimate.

4.4 Results at higher levels of aggregation

The results above pertain to the *class* level of aggregation of COICOPs, which is the lowest level of aggregation in the data. However, the degree of substitutability of goods and services

depends on the granularity in which they are defined. For this reason, we present results for the two higher levels of aggregation, the *group* and *division* classifications, in Table 3.⁴⁵

For reference purposes, Column (I) of the table contains the benchmark results at the class level of aggregation from Table 1. Columns (II) and (III) contain the results at the group and division levels, respectively. Note that the number of COICOPs decreases from 74 at the class level to 10 at the division level.⁴⁶

The first thing to notice in Table 3 is that our point estimate of $\beta_{12,48}$ is increasing in the level of aggregation. This implies a lower elasticity of substitution at higher levels of aggregation. In fact, our point estimate of the implied elasticity for the division level is one, which coincides with Cobb-Douglas preferences.

At higher levels of aggregation the size of the cross-section dimension in the data is, by definition, smaller. Thus the number of observations, listed in Table 3, declines as we move from Column (I) to (III). As a result, the precision with which β is estimated decreases due to a reduction in the degrees of freedom.

The lower precision of the estimates at the higher levels of aggregation is reflected in the wider confidence intervals for the implied elasticity ε . Note, however, that none of the intervals includes elasticities that are substantially higher than 5.

Finally, as for our mapping between $\beta_{12,48}$ and ε at the class level, estimates at higher levels of aggregation also depend on our assumption for α . Figure 4 depicts the estimated elasticity (ε) at the group and division levels for different assumptions for α . It shows that as the labor share increases, the point estimates for the elasticity at the group and division levels of aggregation also increase. As the labor share increases from 2/3, our benchmark level, to 85%, for example, point estimates for the upper-level elasticity of substitution increase from 2.6 to 7.4 at the group level, and from 1 to 2.8 at the division level.⁴⁷

The bottom line is that when one has to choose ε , the appropriate choice depends on the level of aggregation of the consumption aggregate considered. For a high level of aggregation, Cobb-Douglas preferences are consistent with our point estimates. For lower levels of aggregation our results point to an elasticity of substitution of about 3. Taking into account the uncertainty

⁴⁵The aggregation of the data results in potential measurement errors in VAT rates because of within category variation in VAT rates. In addition, estimated coefficients may suffer from attenuation bias. To abate this measurement error, we exclude categories in which the within-group or within-division standard deviation of VAT rates exceeds one percentage point. For more details, please, refer to Appendix A.1.

⁴⁶The Appendix Tables A2 through A4 contain a detailed list of all COICOPs in our dataset.

⁴⁷Unreported 95% confidence intervals for the group and division levels show similar patterns as those at the class level, that is, as the labor share increases, precision of our estimates declines. Assuming the same uniform distribution for α as in our robustness check in Section 4.3 we again obtain point estimates that are very similar to the ones reported in Table 3. The confidence intervals at the group and division levels in that case widen to [1.0, 6.0] and [0.0, 4.4], respectively.

around our estimates we obtain an upper-bound of approximately 5.

5 Implications

The empirical results imply that, depending on the level of aggregation, a reasonable choice of elasticity of substitution for the upper level of nested CES preferences ranges from 1 to 5. Here we discuss two strands of the literature for which this elasticity is relevant and its value has important quantitative implications.

Throughout, it is important to realize that the larger ε , the quicker consumers substitute towards goods with the lowest relative price. In the limit, when $\varepsilon \rightarrow \infty$, the only thing that matters for utility is the lowest price charged across goods. In this case, if relative prices change, but the lowest price remains the same, nothing happens to demand.

Our first example, related to gains from trade, is a case in which relative prices are set efficiently. Our second example is that of the distortion to relative prices due nominal rigidities. These distortions result in a welfare loss that monetary policy could possibly reduce.

Gains from trade The elasticity of substitution between goods and services is of first-order importance for the magnitude of the gains from trade. As [Costinot and Rodríguez-Clare \(2014\)](#) point out, most analyses of the gains from trade in multisector models follow [Dawkins, Srinivasan and Whalley’s \(2001\)](#) “idiot’s law of elasticities” which mandates assuming that the upper-level elasticity of substitution is one (aggregate preferences are Cobb-Douglas), until shown to be otherwise. Our estimates, especially those in Column (III) of Table 3, show that this might not be such an “idiotic” choice after all when one defines the upper-level at a high level of aggregation of goods and services.

In multisector models, the magnitude of the gains from trade depends a lot on the value of the upper-level elasticity of substitution, ε , as trade results in changes to relative price of goods. [Costinot and Rodríguez-Clare \(2014\)](#) show that, in a relatively standard trade model, gains from trade decline in ε . This is because, when there is very little substitution across goods (small ε), autarky implies the absence of certain goods, which, by assumption, cannot easily be substituted. Consequentially, the immediate gain from opening to trade, and the availability of a wider set of goods, are larger the smaller is ε . In particular, [Costinot and Rodríguez-Clare \(2014\)](#) show that gains from opening to trade for $\varepsilon = 1$ tend to be more than double those at $\varepsilon = 5$. However, if globalization continuously lowers relative prices of tradable goods and services, in the long-run this will yield higher welfare gains in the case in which consumers can substitute towards the goods that globalization renders cheaper over time. In that case, the benefits from globalization are increasing in the upper-level of substitution, ε . Moreover,

if globalization results in gains from varieties, as in [Broda and Weinstein \(2006\)](#), those gains from varieties effectively result in relative price declines of goods where these gains are largest. Again, in such a case, the benefits of such gains are higher in case ε is higher.

Distortions from price rigidities Another example where the value of the upper-level elasticity of substitution between expenditure categories matters is in models with nominal rigidities due to price stickiness.⁴⁸ In those models, the inability of producers to adjust their prices distorts relative prices across goods. Such distortions occur because of sector-specific shocks or due to different degrees of price rigidities across goods. This distortion is the one that monetary policy can, partly or completely depending on the model, offset ([Blanchard and Galí, 2007](#)).

The higher the upper-level of substitution, the more demand shifts to the goods and services with distorted low relative prices. Thus, the higher is the upper-level elasticity of substitution the more the allocation of resources in such models is distorted due to nominal rigidities. Everything else equal, when aggregate shocks hit the (multisector) economy and relative prices are distorted, a larger ε results in quicker substitution to relatively cheaper goods, which entails a larger deviation from the flexible price (undistorted) case.⁴⁹

6 Conclusion

Even though nested CES preferences are the most commonly used functional form in multisector macroeconomic models, there are very few estimates of the upper-level elasticity of substitution for this preference specification. Those that do exist either do not cover all expenditure categories over which aggregate preferences are defined, or are based on estimation methods that inherently suffer from endogeneity bias.

In this paper we estimated this upper-level elasticity of substitution using data on all expenditure categories included in the calculation of the HICP. Moreover, our estimate of the elasticity is identified from the long-run response of relative prices to changes in relative VAT rates. Since these VAT rate changes are essentially exogenous, our estimate is not susceptible to endogeneity bias.

Using newly constructed data on monthly inflation and VAT rates for 74 expenditure categories in 25 E.U. countries from 1996 to 2015, we estimate the upper-level elasticity of substi-

⁴⁸For example, in models such as those of [Aoki \(2001\)](#), [Carvalho \(2006\)](#), [Bouakez, Cardia and Ruge-Murcia \(2009\)](#), [Nakamura and Steinsson \(2010\)](#), [Carvalho and Nechio \(2011\)](#), and [Carvalho and Nechio \(2016\)](#), among others.

⁴⁹[Carvalho and Nechio \(2016, Section 3\)](#) provide a thorough discussion on how the elasticities of substitution, particularly the upper-level elasticity, can affect the degree of monetary non-neutrality in alternative multisector economies in which labor (and capital) is sector-specific, firm-specific or economy-wide (as we consider in our model of [Section 2](#)).

tution to be between 1 for the 1-digit level of aggregation of expenditures and 3 for the 3-digit level. Taking into account the uncertainty around these estimates, a value of 5 for the elasticity at the low (3-digit) level of aggregation is the upper bound for the elasticity for which we find support in the data.

Our estimates are meant to provide more informed guidance for the choice of this parameter in quantitative macroeconomic studies. In many applications this parameter has been set to equal 1, such that upper-level preferences are Cobb-Douglas. Our results suggest that this is not an unreasonable choice, as long as this upper-level is defined over 1-digit level expenditure aggregates. At lower levels of aggregation an upper-level elasticity of substitution of 3 is consistent with our estimates.

References

- Allington, Nigel F.B., Paul A. Kattuman, and Florian A. Waldmann. 2005. "One Market, One Money, One Price?" *International Journal of Central Banking*, 1(3).
- Andrade, Philippe, Agnès Bénassy-Quéré, and Martine Carré. 2015. "Competition and pass-through on international markets: Firm level evidence from VAT shocks." Mimeo. Available at <http://www.parisschoolofeconomics.com/benassy-quere-agnes/Papers/VATPassThrough.pdf>.
- Aoki, Kosuke. 2001. "Optimal monetary policy responses to relative-price changes." *Journal of Monetary Economics*, 48(1): 55–80.
- Basu, Susanto, and John G. Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*, 105(2): 249–283.
- Benedek, Dora, Ruud A. de Mooij, and Philippe Wingender. 2015. "Estimating VAT Pass Through." International Monetary Fund IMF Working Papers 15/214.
- Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics*, 25(2): 242–262.
- Berry, Steven T., James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy*, 112(1): 68–105.
- Blanchard, Olivier, and Jordi Galí. 2007. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit and Banking*, 39: 35–65.
- Bouakez, Hafedh, Emanuela Cardia, and Francisco Ruge-Murcia. 2009. "The Transmission of Monetary Policy in a Multi-sector Economy." *International Economic Review*, 50: 1243–1266.
- Broda, Christian, and David E. Weinstein. 2006. "Globalization and the Gains From Variety." *The Quarterly Journal of Economics*, 121(2): 541–585.
- Caliendo, Lorenzo, and Fernando Parro. 2015. "Estimates of the Trade and Welfare Effects of NAFTA." *Review of Economic Studies*, 82(1): 1–44.
- Calvo, Guillermo. 1983. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics*, 12: 383–398.

- Carare, Alina, and Stephan Danninger.** 2008. “Inflation Smoothing and the Modest Effect of VAT in Germany.” International Monetary Fund Working Paper 08/175.
- Carvalho, Carlos.** 2006. “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks.” *Frontiers of Macroeconomics*, 2(1).
- Carvalho, Carlos, and Fernanda Nechio.** 2011. “Aggregation and the PPP Puzzle in a Sticky-Price Model.” *The American Economic Review*, 101(6): 2391–2424.
- Carvalho, Carlos, and Fernanda Nechio.** 2016. “Factor specificity and real rigidities.” *Review of Economic Dynamics*, 22: 208 – 222.
- Cashin, David, and Takashi Unayama.** 2016. “Measuring Intertemporal Substitution in Consumption: Evidence from a VAT Increase in Japan.” *Review of Economics and Statistics*, 98(2): 285–297.
- Costinot, Arnaud, and Andrés Rodríguez-Clare.** 2014. “Chapter 4 - Trade Theory with Numbers: Quantifying the Consequences of Globalization.” In *Handbook of International Economics*. Vol. 4 of *Handbook of International Economics*, , ed. Gita Gopinath, Elhanan Helpman and Kenneth Rogoff, pp. 197–261. Elsevier.
- Dawkins, Christina, T.N. Srinivasan, and John Whalley.** 2001. “Chapter 58 - Calibration.” In *Handbook of Econometrics*. Vol. 5 of *Handbook of Econometrics*, , ed. James J. Heckman and Edward Leamer, pp. 3653–3703. Elsevier.
- Deaton, Angus, and John Muellbauer.** 1980. “An Almost Ideal Demand System.” *The American Economic Review*, 70(3): 312–326.
- Dixit, Avinash K., and Joseph E. Stiglitz.** 1983. “Monopolistic Competition and Optimum Product Diversity.” *American Economic Review*, 67: 297–308.
- Eaton, Jonathan, and Samuel Kortum.** 2002. “Technology, Geography, and Trade.” *Econometrica*, 70(5): 1741–1779.
- European Commission.** 2007. “Employment in Europe.” European Commission, available at ec.europa.eu/social/BlobServlet?docId=3068&langId=en.
- European Commission.** 2015. “VAT Rates Applied in the Member States of the European Union.” European Commission, available at http://ec.europa.eu/taxation_customs/resources/documents/taxation/vat/how_vat_works/rates/vat_rates_en.pdf.

- Eurostat.** 2009. “HICP CT manual.” European Commission, available at http://ec.europa.eu/eurostat/ramon/statmanuals/files/hicp-ct_manual_EN.pdf.
- Eurostat.** 2015. “HICP at constant tax rates.” European Commission, available at http://ec.europa.eu/eurostat/documents/272892/272986/HICP-CT_VAT.pdf.
- Eusepi, Stefano, Bart Hobijn, and Andrea Tambalotti.** 2011. “CONDI: A Cost-of-Nominal-Distortions Index.” *American Economic Journal: Macroeconomics*, 3(3): 53–91.
- Feenstra, Robert C., Philip A. Luck, Maurice Obstfeld, and Katheryn N. Russ.** 2014. “In Search of the Armington Elasticity.” National Bureau of Economic Research NBER Working Paper 20063.
- Fisher, Franklin M.** 1969. “The Existence of Aggregate Production Functions.” *Econometrica*, 37(4): 553–577.
- Gabriel, Peter, and Adam Reiff.** 2010. “Price setting in Hungary: A store-level analysis.” *Managerial and Decision Economics*, 31: 161–176.
- Gautier, Erwan, and Antoine Lalliard.** 2014. “How do VAT changes affect inflation in France?” *Banque de France Quarterly Selection of Articles*, 32.
- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi.** 2014. “Chapter 6 – Growth and Structural Transformation.” In *Handbook of Economic Growth*. Vol. 2, , ed. Philippe Aghion and Steven N. Durlauf, pp. 855–941. Elsevier.
- Hobijn, Bart, Federico Ravenna, and Andrea Tambalotti.** 2006. “Menu Costs at Work: Restaurant Prices and the Introduction of the Euro.” *The Quarterly Journal of Economics*, 121(3): 1103–1131.
- Imbs, Jean, and Isabelle Mejean.** 2015. “Elasticity Optimism.” *American Economic Journal: Macroeconomics*, 7(3): 43–83.
- Jaimovich, Nir.** 2007. “Firm dynamics and markup variations: Implications for sunspot equilibria and endogenous economic fluctuations.” *Journal of Economic Theory*, 137(1): 300–325.
- Jordà, Òscar.** 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *The American Economic Review*, 95(1): 161–182.
- Karadi, Peter, and Adam Reiff.** 2007. “Menu Costs and Inflation Asymmetries Some Micro Data Evidence.” Institute of Economics, Centre for Economic and Regional Studies, Hungarian Academy of Sciences IEHAS Discussion Paper 0706.

- Karadi, Peter, and Adam Reiff.** 2014. “Menu Costs, Aggregate Fluctuations, and Large Shocks.” CEPR Discussion Paper 10138.
- Klenow, Peter J., and Oleksiy Kryvtsov.** 2008. “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *The Quarterly Journal of Economics*, 123(3): 863–904.
- Leslie, Phillip.** 2004. “Price Discrimination in Broadway Theater.” *The RAND Journal of Economics*, 35(3): 520–541.
- Long, John B., and Charles I. Plosser.** 1983. “Real Business Cycles.” *Journal of Political Economy*, 91(1): 39–69.
- Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz.** 1987. “Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment.” *The American Economic Review*, 77(3): 251–277.
- Midrigan, Virgiliu.** 2011. “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations.” *Econometrica*, 79(4): 1139–1180.
- Nakamura, Emi, and Jón Steinsson.** 2010. “Monetary Non-neutrality in a Multisector Menu Cost Model.” *The Quarterly Journal of Economics*, 125(3): 961–1013.
- Nevo, Aviv.** 2001. “Measuring Market Power in the Ready-to-Eat Cereal Industry.” *Econometrica*, 69(2): 307–342.
- Ngai, L. Rachel, and Christopher A. Pissarides.** 2007. “Structural Change in a Multi-sector Model of Growth.” *American Economic Review*, 97(1): 429–443.
- Petrin, Amil.** 2002. “Quantifying the Benefits of New Products: The Case of the Minivan.” *Journal of Political Economy*, 110(4): 705–729.
- Seely, Antony.** 2013. “VAT : the temporary cut in the standard rate.” House of Commons: Business and Transportation Section, available at <http://researchbriefings.files.parliament.uk/documents/SN00701/SN00701.pdf>.
- Simonovska, Ina, and Michael E. Waugh.** 2014a. “The elasticity of trade: Estimates and evidence.” *Journal of International Economics*, 92(1): 34–50.
- Simonovska, Ina, and Michael E. Waugh.** 2014b. “Trade Models, Trade Elasticities, and the Gains from Trade.” National Bureau of Economic Research, Inc NBER Working Papers 20495.

Wilcox, David W. 1992. “The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work.” *The American Economic Review*, 82(4): 922–941.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Wooldridge, Jeffrey M. 2001. *Econometric Analysis of Cross Section and Panel Data*. Vol. 1 of *MIT Press Books*, The MIT Press.

Table 1: Estimates at the class level for $h = 48$ for various model specifications

	(I)	(II)	(III)
β	0.40 (0.04)	0.38 (0.06)	0.47 (0.04)
No. of obs.	293,625	148,270	254,222
Degrees of freedom	268,650	125,702	229,954
Sample:			
Period	1996-2015	1996-2007	1996-2015
No. of COICOPs	74	74	74
Lag inflation	No	No	Yes
Implied elasticity:			
ε	2.98	3.22	2.27
95% conf. interval	[2.1,4.3]	[1.9,5.7]	[1.7,3.1]

Note: Standard errors in parentheses. Implied elasticities calculated using $\alpha = 1/3$. Because ε is a non-linear function of β , we construct the 95% confidence interval for ε by calculating implied elasticities for the boundaries of the symmetric confidence interval for estimated elasticity, β . This results in an asymmetric confidence interval for ε .




Figure 1: Identification strategy




Figure 2: Time series of variation in VAT rate changes




Figure 3: Estimated reduced-form parameter, β , for baseline and unrestricted specifications.

Table 2: Robustness estimates for $h = 48$

	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)
	Benchmark (Table 1, Column I)		Tradability		Clustering by:		Variation within:
	Tradable	Nontradable	Group	Division	Group	Division	
β	0.4 (0.04)	0.36 (0.06)	0.38 (0.07)	0.40 (0.06)	0.40 (0.06)	0.26 (0.07)	0.42 (0.05)
Number of obs.	293,635	182,213	101,151	293,635	293,635	293,635	293,635
Degrees of freedom	268,666	165,288	89,816	94,903	39,293	184,410	236,144
Implied elasticity:	2.98 (0.2)	3.5 (0.3)	3.22 (0.4)	2.98 (0.3)	2.98 (0.3)	5.64 (0.6)	2.74 (0.3)
95% conf. interval	[2.1,4.3]	[2.2,5.9]	[1.8,6.4]	[1.9,5.0]	[1.8,5.1]	[2.9,15.3]	[1.8,4.3]

Note: Standard errors in parentheses. Implied elasticities calculated using $\alpha = 1/3$. Because ε is a non-linear function of β , we construct the 95% confidence interval for ε by calculating implied elasticities for the boundaries of the symmetric confidence interval for estimated elasticity, β . This results in an asymmetric confidence interval for ε . Sample includes data from 1996 to 2015.

Table 3: Estimates for $h = 48$ for various levels of aggregation

	(I) <i>Class</i>	(II) <i>Group</i>	(III) <i>Division</i>
β	0.40 (0.04)	0.43 (0.08)	0.66 (0.17)
No. of obs.	293,635	128,171	16,863
Degrees of freedom	268,666	115,061	11,466
Sample:			
Period	1996-2015	1996-2015	1996-2015
No. of COICOPs	74	36	10
Lag inflation	No	No	No
Implied elasticity:			
ε	2.98	2.61	1.01
95% conf. interval	[2.1,4.3]	[1.4,5.1]	[0.0,4.2]

Note: Standard errors in parentheses. Implied elasticities calculated using $\alpha = 1/3$. Because ε is a non-linear function of β , we construct the 95% confidence interval for ε by calculating implied elasticities for the boundaries of the symmetric confidence interval for estimated elasticity, β . This results in an asymmetric confidence interval for ε .




Figure 4: Estimated elasticity ε for varying labor share $(1 - \alpha)$

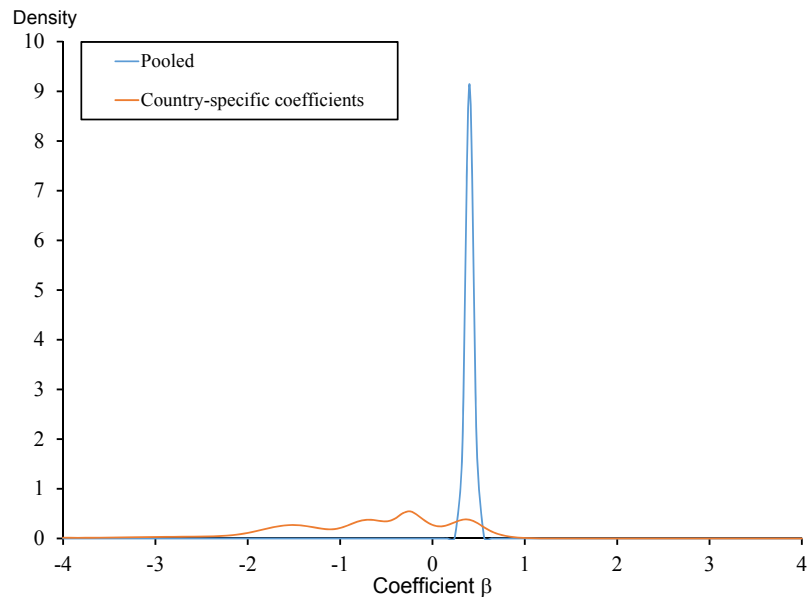


Figure 5: Distributions of pooled and non-pooled coefficients

A Appendix

A.1 Data details

We measure inflation using the change in the logarithm of harmonized consumer price indices (HICP) by consumer expenditure category. Expenditures covered by HICPs are classified in categories/goods called COICOPs. Our data covers the time period from January 1996 to January 2015.

The COICOP classification system consists of three levels of aggregation of categories. In our sample, the top level consists of 10 *divisions*. Our data originally included 12 divisions, from which we dropped all classes (and groups) pertaining divisions 6 (Health) and 10 (Education) because of the non-market price setting in these two sectors. In addition, we do not include two additional divisions that cover spending by non-profit institutions serving households (NPISH) and government consumption. The next level of aggregation is called a *group*. Our sample includes 36 groups. The lowest level of aggregation is a *class*. Our sample includes 74 classes.

Data on VAT rates by COICOP for the 25 European Union countries in our sample are not readily available. We construct them from two administrative sources: [European Commission \(2015\)](#), and [Eurostat \(2015\)](#). These give us information about which VAT rates are applicable to goods and services in a country at a point in time.

To construct our data set, we match all COICOPs for which we have price data for to their corresponding VAT rates at the lowest level of aggregation – the class level. VAT rates for the group- and division-level categories are (for the most part) constructed as weighted averages of the VAT rates in the underlying classes.

To aggregate VAT rates up from the class to the group level, whenever the applied VAT is not directly listed on either [European Commission \(2015\)](#) or [Eurostat \(2015\)](#), we apply a simple average of VAT rates across all classes within each group.

To aggregate up from the group level to the division level, we make use of expenditure data by group and division. In particular, we use Eurostat data on consumer expenditures by COICOP at the group and division levels.⁵⁰ For each country and COICOP, consumer expenditures are available at an annual frequency up to 2013. We match our monthly data set of VAT rates and prices to yearly consumer expenditure shares assuming the latter are unchanged across all months of each year, and using the closest year to fill in missing expenditures data. When matching COICOPs to yearly data on expenditures, for each division at a certain year, we either match expenditure values for all the groups or for none of them, assigning missing

⁵⁰We download the series named “nama-co3-k.” For all countries and categories, units are in millions of 2005 euros.

values in the latter case. We do so to avoid mixing and matching expenditure values from different years. We then compute the share of expenditures in each group over total division expenditures. Finally, whenever the applied VAT at the group level is not directly listed on either [European Commission \(2015\)](#) or [Eurostat \(2015\)](#), we aggregate VAT rates up from the group to the division level by calculating the expenditure-share weighted averages of the VAT rates across groups within each division.

Because this aggregation process potentially introduces noise to our measures of VAT rate, for both the group-level and the division-level estimations, we exclude categories in which the within-category dispersion of VAT rates is larger than one percentage point.⁵¹

Tables [A1](#) to [A4](#) report details of the data that we actually include in our sample after performing the adjustments just described.

Table [A1](#) lists all countries included in the sample. For each country and level of aggregation, it reports the number of 12-month VAT rate changes (between 1996 to 2015), the number of categories (COICOPs) with non-zero VAT rate changes, and the number of non-zero *demeaned* changes in VAT rates that result from the absorption of COICOP-country-month (α_{jcm}) and country-time (γ_{ct}) fixed effects. The latter variation is the one we actually explore in our empirical approach (see equation (24) in the main text document).

Tables [A2](#) to [A4](#) provide analogous information by expenditure category. In particular, they report descriptive statistics of VAT variation for all categories at all levels of aggregation. The first column reports the number of nonzero changes in VAT rates within 12-month for each category. The second column shows the number of countries with nonzero VAT rate changes (out of our sample of 25 countries) for each COICOP.⁵² These two columns report statistics based on simple 12-month changes in VAT rates. The last column of Tables [A2](#) to [A4](#) reports the number of nonzero demeaned-VAT rate changes.

Along with Figure 2 reported in the main text, Tables [A1](#) to [A4](#) suggest that the source of variability in our data does not rely on any specific time, country or category.

A.2 Additional empirical results

For completeness, in this section we report additional results for the group and division levels. In particular, Figures [A1](#) and [A2](#) report the estimated group- and division-levels $\beta_{l,h}$ for $l = 12$ and $h = 12, \dots, 62$, respectively. The figures show that the patterns observed at the class level

⁵¹Estimated coefficients at the group and division levels that include all categories irrespective of the within-level VAT rate dispersion suffer from attenuation bias or are not statistically significant.

⁵²More specifically, the tables report the number of countries with tax changes on a arbitrarily chosen date, December 2013.

are similar at the group and division levels, despite the much wider 95% confidence bands.

Table A5 reports the group-level estimates for β and the implied elasticity ε for the full sample, for the 1996-2007 period, and for the version that includes lagged inflation as an additional control variable. Table A6 performs the analogous exercise at the division level. Column (I) of Tables A5 and A6 replicate the results reported in Table 3 on the main text document. Columns (II) and (III) in both Tables A5 and A6 show that, at these levels of aggregation, because the sample size is much more limited, the precision of our estimates is much reduced, particularly when the sample is further limited by dropping the later part of the time sample (Column (II)). For both estimates at the group and the division levels, further limiting the sample makes the implied elasticity (ε) estimates not statistically different from zero.

Table A1: Variation in VAT rate changes by country

	Class		Group		Division				
	No. of $\Delta\tau_t \neq 0$	COICOPs with $\Delta\tau_t \neq 0$	Demeaned $\Delta\tau_t \neq 0$	No. of $\Delta\tau_t \neq 0$	COICOPs with $\Delta\tau_t \neq 0$	Demeaned $\Delta\tau_t \neq 0$	No. of $\Delta\tau_t \neq 0$	COICOPs with $\Delta\tau_t \neq 0$	Demeaned $\Delta\tau_t \neq 0$
Austria	0	0	0	0	0	0	192	1	868
Belgium	24	2	15841	12	1	7161	0	0	0
Cyprus	3120	46	14544	1496	24	6348	299	3	963
Czech Republic	2700	67	15841	1020	27	6727	156	4	868
Denmark	0	0	0	0	0	0	0	0	0
Estonia	828	65	15745	372	28	7547	60	5	0
Finland	2052	65	16058	654	25	6293	93	3	651
France	1296	53	15679	552	23	6510	60	3	855
Germany	1320	54	15841	624	26	6944	96	4	813
Greece	2014	64	15841	690	24	6293	90	3	606
Hungary	2438	66	15841	972	27	6998	187	5	936
Ireland	2568	53	15841	948	17	4774	184	2	534
Italy	1584	42	15841	792	22	6293	108	3	633
Latvia	2381	61	15275	996	27	7049	144	4	868
Lithuania	1512	65	15902	696	29	7161	120	5	765
Luxembourg	51	50	1387	18	18	475	193	1	434
Malta	588	36	15624	264	20	6076	12	1	0
Netherlands	1056	39	15624	504	19	6293	48	2	651
Poland	948	69	15841	300	24	6076	153	3	627
Portugal	2862	67	16058	1332	27	6727	178	3	709
Slovakia	2472	65	15793	1080	30	7595	228	6	1338
Slovenia	1644	65	12985	540	22	5049	72	3	525
Spain	1344	49	15841	576	24	6293	284	4	1056
Sweden	0	0	0	0	0	0	0	0	0
United Kingdom	1848	47	15841	948	25	6727	144	4	204

Table A2: Variation in VAT rate changes by expenditure category

Expenditure Categories:		No. of $\Delta\tau_t \neq 0$ within 12-months	No. of countries with $\Delta\tau_t \neq 0$	No. of demeaned $\Delta\tau_t \neq 0$ changes
1	Food and non-alcoholic beverages	618	11	2835
1.1	Food	295	10	4263
1.11	Bread and cereals	331	11	4480
1.12	Meat	331	11	4480
1.13	Fish and seafood	331	11	4480
1.14	Milk, cheese and eggs	331	11	4480
1.15	Oils and fats	331	11	4480
1.16	Fruit	331	11	4480
1.17	Vegetables	331	11	4480
1.18	Sugar, jam, honey, chocolate and confectionery	331	11	4480
1.19	Food products n.e.c.	331	11	4480
1.2	Non-alcoholic beverages	319	11	2961
1.21	Coffee, tea and cocoa	343	12	4480
1.22	Mineral waters, soft drinks, fruit and vegetable juices	523	18	4480
2	Alcoholic beverages, tobacco and narcotics	823	20	3909
2.1	Alcoholic beverages	572	19	4298
2.11	Spirits	627	21	4534
2.12	Wine	591	21	4534
2.13	Beer	627	21	4534
2.2	Tobacco	627	21	4534
3	Clothing and footwear	542	13	2691
3.1	Clothing	442	13	3038
3.11	Clothing materials	627	17	4534
3.12	Garments	627	21	4534
3.13	Other articles of clothing and clothing accessories	627	21	4534
3.14	Cleaning, repair and hire of clothing	561	19	4534
3.2	Footwear	627	21	4534
4	Housing, water, electricity, gas and other fuels	0	0	0
4.1	Actual rentals for housing	0	0	4576
4.3	Maintenance and repair of the dwelling	317	13	2840
4.31	Materials for the maintenance and repair of the dwelling	627	21	4534
4.32	Services for the maintenance and repair of the dwelling	425	18	4480
4.4	Water supply and miscellaneous services relating to the dwelling	168	6	1302
4.41	Water supply	288	11	4480
4.42	Refuse collection	396	14	4480
4.43	Sewerage collection	520	18	4480
4.44	Other services relating to the dwelling n.e.c.	627	19	4534
4.5	Electricity, gas and other fuels	346	12	2779
4.51	Electricity	537	20	4450
4.52	Gas	537	19	4450
4.53	Liquid fuels	555	16	4534
4.54	Solid fuels	572	16	4388
4.55	Heat energy	553	12	4482

Table A3: Variation in VAT rate changes by expenditure category (cont.)

Expenditure Categories:		No. of $\Delta\tau_t \neq 0$ within 12-months	No. of countries with $\Delta\tau_t \neq 0$	No. of deemed $\Delta\tau_t \neq 0$ changes
5	Furnishings, household equipment and routine household maintenance	470	15	3007
5.1	Furniture and furnishings, carpets and other floor coverings	615	20	4317
5.11	Furniture and furnishings	615	20	4534
5.12	Carpets and other floor coverings	627	21	4534
5.13	Repair of furniture, furnishings and floor coverings	627	10	4534
5.2	Household textiles	627	21	4534
5.3	Household appliances	627	21	4534
5.31	Major and small household appliances	627	21	4534
5.33	Repair of household appliances	627	21	4534
5.4	Glassware, tableware and household utensils	627	21	4534
5.5	Tools and equipment for house and garden	627	21	4534
5.6	Goods and services for routine household maintenance	482	16	3689
5.61	Non-durable household goods	627	21	4534
5.62	Domestic services and household services	591	21	4534
7	Transport	443	6	1179
7.1	Purchase of vehicles	627	21	4534
7.11	Motor cars	627	21	4534
7.12	Motor cycles, bicycles and animal drawn vehicles	627	20	4534
7.2	Operation of personal transport equipment	442	13	3038
7.21	Spare parts and accessories for personal transport equipment	627	21	4534
7.22	Fuels and lubricants for personal transport equipment	627	21	4534
7.23	Maintenance and repair of personal transport equipment	561	19	4534
7.24	Other services in respect of personal transport equipment	627	21	4534
7.3	Transport services	144	5	1085
7.31	Passenger transport by railway	347	14	4016
7.32	Passenger transport by road	347	15	4480
7.33	Passenger transport by air	375	14	4576
7.34	Passenger transport by sea and inland waterway	375	9	4534
7.35	Combined passenger transport	48	1	651
7.36	Other purchased transport services	603	13	4317
8	Communications	0	0	0
8.1	Postal services	0	0	4359
8.2	Telephone and telefax equipment	567	16	4317
8.3	Telephone and telefax services	567	16	4317

Table A4: Variation in VAT rate changes by expenditure category (cont2.)

Expenditure Categories:		No. of $\Delta\tau_t \neq 0$ within 12-months	No. of countries with $\Delta\tau_t \neq 0$	No. of demeaned $\Delta\tau_t \neq 0$ changes
9	Recreation and culture	12	0	36
9.1	Audio-visual, photographic and information processing equipment	627	21	4534
9.1.1	Equipment for sound and picture	627	21	4534
9.1.2	Photographic and cinematographic equipment and optical instruments	627	20	4534
9.1.3	Information processing equipment	627	21	4534
9.1.4	Recording media	627	21	4534
9.1.5	Repair of audio-visual, photographic and information processing equipment	627	20	4534
9.2	Other major durables for recreation and culture	627	19	4534
9.2.1	Major durables for indoor and outdoor recreation	627	19	4534
9.2.3	Maintenance and repair of other major durables for recreation and culture	627	4	4534
9.3	Other recreational items and equipment, gardens and pets	603	20	4317
9.3.1	Games, toys and hobbies	627	21	4534
9.3.2	Equipment for sport, camping and open-air recreation	627	21	4534
9.3.3	Gardens, plants and flowers	627	21	4534
9.3.4	Pets and related products; veterinary and other services for pets	627	21	4534
9.4	Recreational and cultural services	382	11	2821
9.4.1	Recreational and sporting services	590	19	4534
9.4.2	Cultural services	573	18	4534
9.5	Newspapers, books and stationery	12	1	217
9.5.1	Books	258	11	4248
9.5.2	Newspapers and periodicals	306	13	4296
9.5.3	Miscellaneous printed matter; stationery and drawing materials	627	21	4534
9.6	Package holidays	627	21	4534
11	Restaurants and hotels	193	7	1247
11.1	Catering services	418	15	3430
11.1.1	Restaurants, cafs and the like	490	18	4318
11.1.2	Canteens	552	18	4534
11.2	Accommodation services	406	16	4157
12	Miscellaneous goods and services	0	0	0
12.1	Personal care	406	13	3152
12.1.1	Hairdressing salons and personal grooming establishments	527	19	4480
12.1.2	Electrical appliances and articles for personal care	627	21	4534
12.3	Personal effects n.e.c.	627	21	4534
12.3.1	Jewellery, clocks and watches	627	21	4534
12.3.2	Other personal effects	627	21	4534
12.4	Social protection	347	12	4576
12.5	Insurance	12	1	4576
12.5.2	Insurance connected with the dwelling	12	1	4576
12.5.3	Insurance connected with health	12	0	4576
12.5.4	Insurance connected with transport	12	1	4576
12.5.5	Other insurance	12	0	4576
12.6	Financial services n.e.c.	0	0	4576
12.7	Other services n.e.c.	627	21	4534

Table A5: Estimates for $h = 48$ for various model specifications at the group level

	(I)	(II)	(III)
β	0.43 (0.08)	-0.15 (0.11)	0.66 (0.08)
No. of obs.	128,171	64,703	110,858
Degrees of freedom	115,061	53,828	98,336
Sample:			
Period	1996-2015	1996-2007	1996-2015
No. of COICOPs	36	36	36
Lag inflation	No	No	Yes
Implied elasticity:			
ε	2.61	-15.62	1.03
95% conf. interval	[1.4,5.1]	[31.2,-7.7]	[0.5,1.9]

Note: Standard errors in parentheses. Implied elasticities calculated using $\alpha = 1/3$. Because ε is a non-linear function of β , we construct the 95% confidence interval for ε by calculating implied elasticities for the boundaries of the symmetric confidence interval for estimated elasticity, β . This results in an asymmetric confidence interval for ε .

Table A6: Estimates for $h = 48$ for various model specifications at the division level

	(I)	(II)	(III)
β	0.66 (0.17)	0.36 (0.21)	0.78 (0.20)
No. of obs.	16,863	8,911	14,515
Degrees of freedom	11,466	5,664	9,718
Sample:			
Period	1996-2015	1996-2007	1996-2015
No. of COICOPs	36	36	36
Lag inflation	No	No	Yes
Implied elasticity:			
ε	1.01	3.56	0.56
95% conf. interval	[0.0,4.2]	[-45.7,0.6]	[-0.3,3.0]

Note: Standard errors in parentheses. Implied elasticities calculated using $\alpha = 1/3$. Because ε is a non-linear function of β , we construct the 95% confidence interval for ε by calculating implied elasticities for the boundaries of the symmetric confidence interval for estimated elasticity, β . This results in an asymmetric confidence interval for ε .




Figure A1: Estimated reduced-form parameter, β , for baseline and unrestricted specifications at the group level.

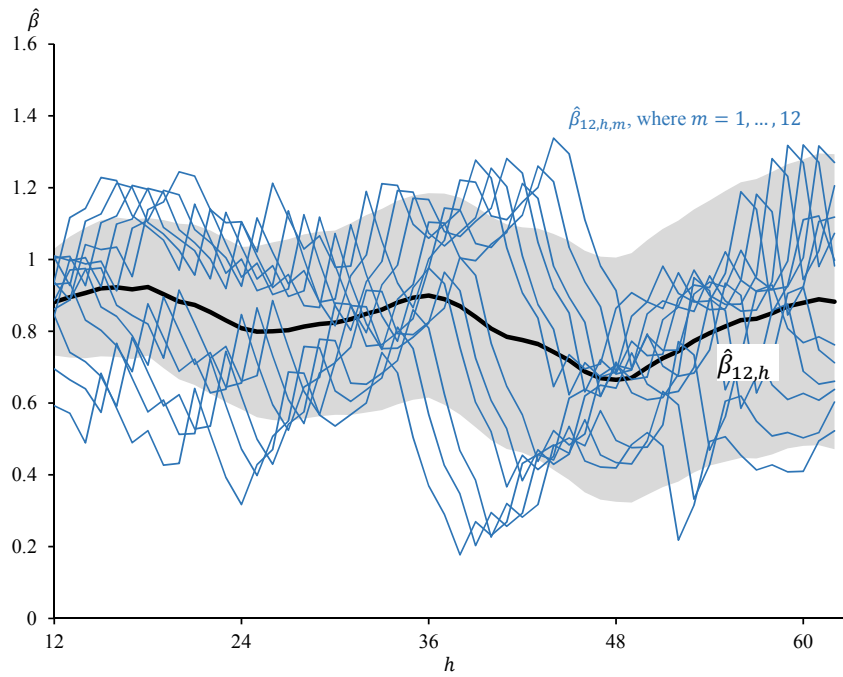


Figure A2: Estimated reduced-form parameter, β , for baseline and unrestricted specifications at the division level.