

What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes

Youssef Benzarti

University of California, Santa Barbara, and National Bureau of Economic Research

Dorian Carloni

Congressional Budget Office

Jarkko Harju

VATT Institute for Economic Research

Tuomas Kosonen

Labour Institute for Economic Research

This paper provides evidence that prices respond significantly more strongly to increases than to decreases in value-added taxes (VATs). First, using two plausibly exogenous VAT changes, we show that prices respond twice as much to VAT increases as to VAT decreases. Second,

Previously circulated as “Do Prices Respond Differently to Increases and Decreases in Consumption Taxes?” The views expressed in this paper are the authors’ and should not be interpreted as those of the Congressional Budget Office, Labour Institute for Economic Research, or VATT. We thank the editor (Magne Mogstad), six anonymous referees, Alan Auerbach, Raj Chetty, Stefano DellaVigna, Ignacio Esponda, Itzik Fadlon, Naomi Feldman, Amy Finkelstein, Xavier Gabaix, Al Harberger, Jim Hines, Hilary Hoynes, Henrik Kleven, Wojciech Kopczuk, Peter Kuhn, Camille Landais, Maurizio Mazzocco, Pascal Michailat, Ryan Oprea, Heather Royer, Emmanuel Saez, Joel Slemrod, Alisa Tazhitdinova, Gonzalo Vazquez-Bare, and numerous seminar participants for their helpful suggestions and comments. Tero Hokkanen, Jen Kades, and Carolina Kansikas provided outstanding research assistance. Harju gratefully acknowledges financial support from the Foundation for Economic Education. Data are provided as supplementary material online.

Electronically published October 26, 2020

[*Journal of Political Economy*, 2020, vol. 128, no. 12]

© 2020 by The University of Chicago. All rights reserved. 0022-3808/2020/12812-0002\$10.00

we show that this asymmetry results in higher equilibrium profits and markups. Third, we find that firms operating with low profit margins are particularly likely to respond asymmetrically to VAT changes. Fourth, these asymmetric price effects persist several years after VAT changes take place. Fifth, using all VAT changes in the European Union from 1996 to 2015, we find similar levels of asymmetry.

I. Introduction

Value-added taxes (VATs) influence a large share of the world's economies, having been adopted in some form by all OECD (Organisation for Economic Co-operation and Development) countries except for the United States. In the European Union, VATs raise 30% of total tax revenue or 12% of GDP, which amounts to the largest source of government revenue. It is therefore important, for both economic and policy reasons, to understand the mechanisms underlying the incidence of VATs. In standard tax incidence models, if the VAT rate is first cut and then increased to its original level and if everything else is held constant, the VAT increase and decrease pass-through rates should be identical, no matter what functional forms or market structure models are assumed. In this paper, we provide empirical evidence calling this fundamental incidence symmetry into question: in our data, responses to VAT increases are substantially larger than responses to VAT decreases and this asymmetry is persistent over time.

This finding is relevant for policy because pass-through estimates used to inform policy typically do not differentiate between VAT increases and decreases and are therefore likely to severely overstate the price effects of VAT decreases and understate the effects of increases. Indeed, according to our estimates, failing to account for the asymmetry we document would lead to overestimates of the pass-through of VAT decreases by a factor of 3. This raises serious fiscal policy concerns, since policy makers have recently successfully advocated cutting VAT rates in an effort to stimulate demand via exactly such overestimated price effects.¹ Moreover, the asymmetric responses we document suggest that using temporary VAT cuts to offer relief to consumers may have the opposite of their intended effect: in the long run they may result in higher equilibrium prices once the VAT cut is repealed, benefiting firm owners at the expense of consumers.

We document this asymmetry using two different approaches. First, we examine the effects of two VAT changes that were plausibly exogenous to the underlying economic conditions. In particular, we study the effects of

¹ Such policies are, moreover, expensive: the United Kingdom's 2009 temporary VAT cut cost an estimated £12.4 billion (see Chirakijja et al. 2009), and France's 2009 sit-down restaurant VAT cut, which is still in place, cost 3 billion euros in 2010 alone (see Benzarti and Carloni 2019).

a 14 percentage point decrease in the VAT rate applied to Finnish hair-dressing services in January 2007 and a subsequent 14 percentage point increase in the same sector in January 2012. We document—using European Commission Council Directives—that the two changes were part of a VAT experimentation program, and therefore that their timing and the choice of sector are plausibly exogenous to underlying economic conditions. Using micro price and corporate tax data, we compare hairdressing services to a control group consisting of beauty salons (which were unaffected by the VAT changes) and report five main results.

First, we find that prices respond twice as much to the 14 percentage point VAT increase as to the 14 percentage point VAT decrease. Second, we find that this asymmetry persists several years after the VAT cut is repealed, suggesting that equilibrium prices depend on the history of tax changes. Third, we find that the asymmetric pass-through is reflected in both markups and profits: both respond asymmetrically to VAT changes and end up at a higher equilibrium level once the VAT cut is repealed. Fourth, we uncover an additional layer of asymmetry: the underlying distribution of price changes following the VAT increase is substantially different from that of the VAT decrease. Fifth, we find that firms operating with low profit margins are particularly likely to respond asymmetrically to the VAT changes.

While the Finnish VAT experiment has several advantages including a large tax change and the availability of rich firm and price data, it suffers from two main shortcomings. First, it is difficult to draw statistical inferences on the magnitude of the asymmetry if one only observes two VAT changes. Second and relatedly, the two VAT changes do not occur at the same time, raising the potential concern that the difference in pass-through is driven by (for example) a change in underlying economic conditions.

We address these shortcomings by turning to a second level of analysis that uses all VAT changes that occurred in the member states of the European Union from 1996 to 2015. Because of the large number of reforms, we can precisely estimate the magnitude of the asymmetry while controlling for several key characteristics of the reforms. We find, using fixed-effects regressions, pass-through asymmetries that are similar in magnitude to asymmetries observed in the Finnish experiment. Unlike the Finnish experiment, these VAT reforms were initiated by the EU member states and their timing, size, the commodity they affect, et cetera, are therefore not random. We address this concern in several ways. First, we show that there are no significant pretrends in prices prior to the reforms. Second, we use matching estimators to match VAT increases and decreases over several key characteristics of the reforms, estimate pass-through rates on the subset of matched reforms, and find similar levels of asymmetric pass-through. This mitigates the concern that VAT increases might be intrinsically different from VAT decreases, which could affect the pass-through of

the tax to prices.² Third, we regress the timing of reforms for both VAT increases and decreases on country-specific economic indicators, including GDP and the unemployment rate. We find no correlation between the timing of VAT changes and the economic conditions leading up to the reforms, which further mitigates concerns that VAT changes are endogenous to economic conditions.

Our paper makes two main contributions. First, we provide evidence that asymmetric responses to VAT changes arise in a large set of countries and for a broad set of commodities. Our results therefore call into question the empirical relevance (at least for one important class of taxes) of an essential part of standard, static tax incidence theory. This in turn suggests that moving beyond static incidence theory and explicitly introducing dynamics when assessing the effects of taxation (possibly using approaches in the spirit of Golosov, Troshkin, and Tsyvinski [2011]) may be crucial for forming accurate welfare conclusions and setting effective policy.

Second, we contribute to a small but growing literature on VATs, which are somewhat understudied given their importance and widespread use.³ The policy and research relevance of the VAT is highlighted in the Mirrlees Review (Mirrlees 2010), which offers key tax reform proposals, with the VAT being prominently featured. The proposals related to the VAT in that review rely on the key assumption that VAT incidence falls fully on consumers—an assumption made by default and, as Atkinson (2012) emphasizes, mostly because of a lack of empirical evidence on that point. Since the publication of the Mirrlees Review, researchers have made significant progress in measuring the incidence of VATs. Gaarder (2018), for example, uses a compelling empirical design to estimate the incidence and distributional effects of a VAT cut on food items in Norway. Our contribution relative to this important literature is to contrast the incidence of VAT increases and decreases. Furthermore, we contribute to our general understanding of the effect of VATs on the economy, following in the same line as papers such as Feldstein and Krugman (1990), Hines and Desai (2005), Kosonen (2015), Pomeranz (2015), Benzarti and Carloni (2019), Naritomi (2019), and Benedek et al. (2020).⁴

² Our matching approach also allows us to construct a sample of VAT changes that is well suited for testing the theoretical predictions we derive in sec. II.

³ A Proquest search of the expression “Value-Added Tax” returns 17,979 scholarly peer-reviewed articles, while “Income Tax” returns 140,408 such articles.

⁴ Notably, Benedek et al. (2020) estimate the pass-through of VATs to prices using sources of data similar to those that we use. While we focus on providing evidence that prices respond asymmetrically to variation in VAT rates, Benedek et al. (2020) estimate the pass-through of VATs. There are also some significant differences in the two approaches, as we use a different empirical framework with a larger set of commodities, countries, and years.

We also contribute to a literature in public finance that estimates tax incidence.⁵ Our paper is the first to provide systematic evidence on the asymmetric pass-through of taxes and to show that prices consistently respond more to increases than to decreases in tax rates. Our paper is related to Carbonnier (2008), but our findings are different. While we show that prices respond systematically more to VAT increases than to decreases, Carbonnier (2008) finds that prices in some industries respond more to VAT increases, while in others they respond more to VAT decreases.⁶ These differences in findings could be due to two limitations of Carbonnier (2008) that are significantly eased in our study: (1) that study considers only 11 commodities in France, while we consider all commodities across the European Union, and, relatedly, (2) it only uses two small VAT changes (a 2 percentage point VAT increase and a 1 percentage point VAT decrease), while our sources of variation are substantially richer and larger. Our results also contrast with those of Doyle and Samphantharak (2008), who find symmetric responses of prices to a 120 day temporary moratorium on a 5% gasoline tax in 2000. There are two possible explanations for the symmetric response found in Doyle and Samphantharak (2008). First, the moratorium was implemented by the governor of Indiana during an election year because he was concerned about the effect of soaring gasoline prices on his reelection. For this reason, gasoline retailers were likely to be under both scrutiny and pressure to reduce prices. Second, because the moratorium lasted only 120 days, asymmetric price changes would have been relatively easy to detect and could have resulted in substantial consumer antagonism.

Our findings are also related to a literature in industrial organization that tests for the asymmetric pass-through of input costs. The fundamental difference between our paper and that literature is that, there, prices tend to exhibit a timing asymmetry when responding to cuts versus increases in input costs and a convergence to symmetry over time. The asymmetry lasts for 1 month in Borenstein, Cameron, and Gilbert (1997), for example, and 3 to 5 months in Peltzman (2000). By contrast, we observe that prices respond immediately to VAT changes and we find no evidence of convergence over time. Further, VAT changes are likely to be a better source of identification than input cost changes. Peltzman (2000), for example, notes that because input costs are not observable (unlike VAT changes), one may estimate spurious asymmetric pass-through rates of input costs.

⁵ Kotlikoff and Summers (1987) and Fullerton and Metcalf (2002) provide a survey of the tax incidence literature.

⁶ Politi and Mattos (2011) and Karadi and Reiff (2019) are two other papers that consider asymmetric responses of prices to VAT reforms. They suffer from some of the same drawbacks as Carbonnier (2008).

II. Theoretical Framework

Empirically, it is important to distinguish between asymmetries that are due to nonconstant pass-through functions and “true asymmetries” where the pass-through rate depends on the direction of the tax change. Standard static incidence models can be consistent with nonconstant pass-through functions in which the pass-through rate depends on the initial level of the tax. Instead, in this paper, we are interested in empirically testing whether pass-through depends on the direction of the tax change, a finding that would be inconsistent with standard static incidence theory. This section provides a framework for disentangling these two effects.

Suppose the government levies an ad valorem tax τ , and let ρ denote the marginal pass-through to prices—the response of prices to a very small change in τ . Further, define the average pass-through of a large change in the ad valorem tax rate T as

$$\kappa(T) = \frac{1}{T} \int_0^T \rho(\tau) d\tau. \quad (1)$$

As long as ρ is assumed continuous, the marginal pass-through rate (i.e., of infinitesimal tax changes) cannot be asymmetric but the average pass-through rate, $\kappa(T)$, of a large tax rate can. We will focus therefore on the properties of $\kappa(T)$ in what follows.

The ideal experiment for testing whether pass-through rates depend on the direction of the tax change would rely on a test in which the VAT rate is first cut and then increased to its original level and everything else is held constant. We use two main sources of variation to mimic this ideal experiment. The first one is a 14 percentage point VAT decrease and subsequent 14 percentage point VAT increase that affect hairdressers in Finland. Here, the pass-through function is integrated over the same range in both the decrease and increase, so the asymmetry cannot be the product of a nonconstant pass-through function, at least under the assumption that nothing else about the market—including other determinants of prices—changes between the two reforms. The second test uses all VAT changes in the European Union from 1996 to 2015. This test helps alleviate the concern that other things may be changing across the two Finnish VAT changes, as explained below.

In the Finnish experiment, the VAT rate is cut starting from a baseline rate and later is increased by the same magnitude. Using the preceding notation, suppose that the VAT rate is reduced by $T \leq t$ from an initial VAT rate of t and then increased back to its original level t . In this case, we would be comparing the average pass-through of a VAT cut of size T , $(1/T) \int_t^{t-T} \rho(\tau) d\tau$, from a baseline rate of t , to the pass-through of a VAT increase of size T , $(1/T) \int_{t-T}^t \rho(\tau) d\tau$, from a baseline VAT rate of $t - T$. By definition, and for any T , t , and $\rho(\cdot)$,

$$\frac{1}{-T} \int_t^{t-T} \rho(\tau) d\tau = \frac{1}{T} \int_{t-T}^t \rho(\tau) d\tau.$$

Because of this, asymmetric pass-through is inconsistent with standard incidence theory predictions for this type of VAT change, for any function $\rho(\cdot)$ and regardless of the degree of competition or the nature of demand and supply functions.

Intuitively, since the VAT rate is first cut and then increased to its original level, the pass-through function is integrated over the same interval and, therefore, standard static incidence models predict that the VAT increase and decrease pass-through rates should be identical. Consequently, any standard incidence model predicts that if a tax rate is decreased and then increased back to its original level, the postincrease equilibrium price should match the predecrease equilibrium price. We investigate this prediction using the Finnish reforms in section IV.

In contrast to the Finnish case, the European VAT changes affect different commodities in different countries and, in general, have different underlying characteristics. Therefore, a direct comparison of their pass-through rates might be confounded by underlying differences in these characteristics. In section V.B we attempt to mitigate these concerns in standard ways by controlling for fixed effects and other observables. In addition, in section V.C, we also conduct a matching exercise that leverages the theoretical observation made above. We match VAT increases and decreases on key characteristics so as to mirror the Finnish experiment. Specifically, preincrease tax rates are matched to postdecrease rates and vice versa in order to make sure that VAT increases and decreases, in the matched sample, average over similar values of the tax rate. This provides a test of standard incidence theory since it would not, with any degree of competition and for any demand and supply functions, generate asymmetries between increases and decreases.⁷

III. Data and Institutional Background

A. Value-Added Taxes

VATs are taxes applied to the value added to transacted goods and services and are included in consumer prices in the European Union. Firms remit the VAT that they collect from consumers to the government and claim credits for the VAT they pay on input costs, ensuring that only the value added is taxed. Final consumers, who are the last component of

⁷ We are very grateful to anonymous referee 6 for suggesting this empirical test and, more generally, for all the helpful and constructive feedback the referee gave us on this section.

the chain of transactions, cannot claim any tax credit and, therefore, pay the tax on the entire value of final goods purchased.

Member countries of the European Union generally have several VAT rates in place, including a standard rate that applies to the majority of commodities, a reduced rate for basic necessities such as food, heating, and passenger transport, and exemptions or zero rates for some commodities.

B. Finnish Hairdressing Sector VAT Reforms

1. Institutional Background

While the European Commission restricts excessive VAT changes to avoid VAT competition, it allows member states to experiment with reduced VAT rates for a small sample of labor-intensive services, with the explicit goal of analyzing the incidence of VATs on prices and employment. The European Commission (1999, 2006) established and explicitly listed the full set of services over which countries are allowed to conduct these experiments. While the list includes hairdressing services, it excludes other very similar services such as beauty salons. This makes hairdressing services a natural treatment group, and beauty salons a plausible control group. Finland took part in the second wave of the experimentation program, which was set to start in January 2007 (Council Directive 2006/112/EC). It was determined in November 2006 that the rate would subsequently revert to its original level. This resulted in a reduction in the VAT rate on hairdressing services from 22% to 8% in January 2007 and a subsequent increase from 9% to 23% in January 2012.⁸ Because the timing, magnitude, and commodities affected by this reform were set by the European Commission, the reforms are plausibly exogenous to economic conditions.

Importantly, there are no special characteristics of the hairdressing sector in Finland that are likely to raise external validity concerns. For example, there are no specific business or licensing requirements imposed on hairdressers that could create barriers to entry. Similarly, the sector does not benefit from any particular status relative to other sectors in the Finnish economy.

Kosonen (2015) previously analyzed the first leg of the reform and part of our analysis replicates some of these results. Specifically, like us, Kosonen (2015) considers the pass-through distribution of the VAT cut (fig. 4A) and estimates the effect of the VAT cut on profits (years 2000 to 2009 in fig. 5A). However, though Kosonen (2015) also documents the effect of the VAT cut on costs, it does not break the effects down by fixed versus

⁸ The reduced and standard VAT rates were both increased by 1 percentage point in July 2010.

variable costs or discusses the possibility of asymmetric pass-through since it does not consider any evidence related to the VAT increase.

2. Data

We use price data collected by surveyors from a random sample of the full population of hairdressers before and after each VAT change. Prices for nine types of services were collected: short-hair haircuts, long-hair haircuts, children's haircuts, complicated haircuts, short-hair permanent waves (perms), long-hair permanent waves, short-hair coloring, long-hair coloring, and complicated coloring. The prices collected are the "menu" prices rather than transaction prices, but we also have information on whether coupons or discounts are offered in each particular location. The data set contains 2,822 price observations for the VAT decrease originating from 427 firms and 2,106 price observations for the VAT increase stemming from 347 firms. We also use micro and aggregate price data from Statistics Finland for haircuts, other hairdressing services, and beauty salons to analyze the long-term effects of the reforms.

We supplement the price data with corporate tax data covering the entire population of firms in Finland. The data are annual and contain information on every line of profits and losses, thus allowing us to observe turnover, fixed and variable costs separately, as well as employee counts. Table B1 shows summary statistics for hairdressers and beauty salons.

C. European VAT Reforms

1. Institutional Background

We examine all VAT changes in the European Union from 1996 to 2015. There are three types of VAT changes in this sample of reforms: (1) changes that affect standard VAT rates that apply to most commodities in the economy, (2) changes that affect reduced VAT rates that apply to commodities that are generally considered necessities, and (3) sector-specific VAT changes.

2. Price Data

We use price data from Eurostat's *Harmonised Indices of Consumer Prices*. The data set contains monthly non-seasonally adjusted information on commodity prices across European countries for the period 1996–2015. Commodities are categorized into Classification of Individual Consumption According to Purpose (COICOP) groups.⁹ These data represent the single most reliable source of information on prices across EU countries. They do not contain information on the prices of intermediate goods.

⁹ See tables E5 and E6 for a list of COICOP groups.

3. Historical VAT Rates

Information on VAT rates by commodity and country is provided directly by the European Commission (EC) in its annual report: *VAT Rates Applied in the Member States of the European Community*. The report contains detailed information on the VAT rate applied to each commodity in each European country, as well as the exact date of the VAT reforms. It covers all commodities subject to VATs.

Because the reports only contain information on current EU members, we consider the following countries: (1) since 1996: Austria, Belgium, Denmark, France, Finland, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain, Sweden, and the United Kingdom; (2) since 2004: Cyprus, the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Slovakia, and Slovenia; (3) since 2007: Bulgaria and Romania. We exclude Croatia because it became a member of the European Union only in 2013.

We drop the education sector from our sample because, here, for-profit institutions are subject to VATs whereas not-for-profit institutions are exempt. The majority of institutions are not-for-profit and are therefore unaffected by the VAT changes, but we cannot differentiate for-profit from not-for-profit institutions in the price data set. We also drop clothing and footwear, as prices in these sectors exhibit strong seasonality, with most sales occurring in January, which is also the month in which most VAT changes occur.

Figures B1 and B2 plot the distribution of VAT increases and decreases by commodity, country, economic conditions (unemployment rate and GDP growth), size, time of the VAT changes, and pre- and postreform VAT rates. Overall, in our sample of VAT changes, 28% are VAT decreases and 83% are economy-wide VAT changes.¹⁰

IV. Finnish Hairdressing Reforms

A. *Price Response to the VAT Changes and Long-Term Persistence*

Figure 1 uses time series from Statistics Finland from January 2005 to November 2015 to show the evolution of hairdressing and beauty salon prices. Prior to the January 2007 reform, the VAT rates for hairdressing services and beauty salons were equal. In January 2007, the VAT for hairdressing services was decreased by 14 percentage points and in January 2012, the rate increased by 14 percentage points; VAT rates for beauty salons, by contrast, were held fixed throughout.

¹⁰ We discuss some of the institutional reasons that explain why there are more VAT increases than decreases in app. B.

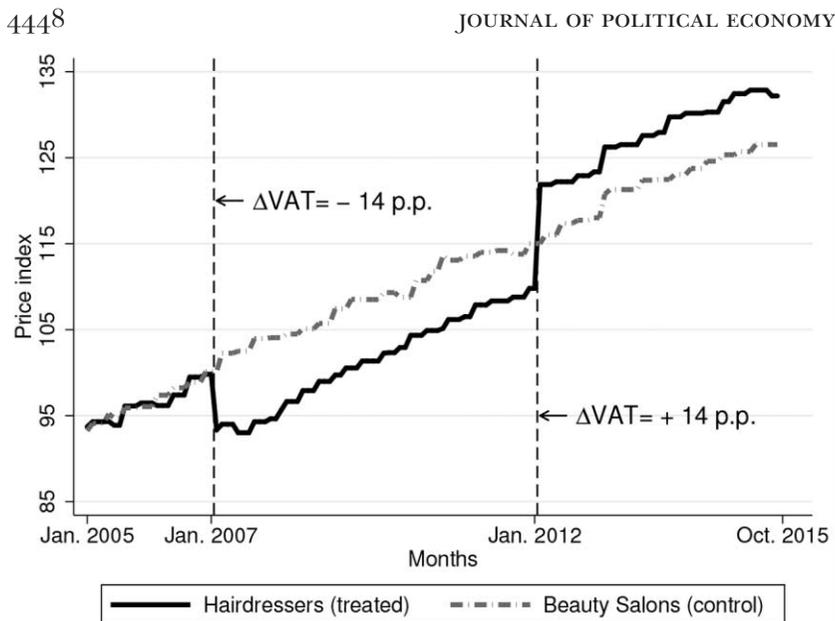


FIG. 1.—Finnish hairdressing sector VAT reforms. This figure shows the price of hairdressing services and beauty salons before and after the 14 percentage point hairdressing services VAT cut in January 2007 and the 14 percentage point VAT hairdressing services hike in January 2012.

Four main empirical patterns emerge from figure 1. First, beauty salons seem to be a natural control group for hairdressing services: prices follow parallel trends throughout the entire prereform period. Second, the largest response of hairdressing prices to both the VAT decrease and increase occurs during the first month following the VAT change. Third, the pass-through rate for the VAT increase is approximately double that of the VAT decrease.¹¹ Fourth, once the VAT rate applied to Finnish hairdressers is increased back to its original level, prices remain higher than for the control group 3.5 years later in spite of the VAT rates being equal for both groups. In section V.D, we provide evidence that asymmetric pass-through is likewise persistent in other markets and countries. This suggests that the market equilibrium depends on the history of tax changes.

B. Pass-Through Distributions

1. Short-Run Pass-Through Distributions

We use micro-level price data to plot the short-run distribution of pass-through rates. We compute the pass-through rate by taking the log

¹¹ We also estimate the pass-through of the VAT increase and decrease to prices separately for each service offered by hairdressers, controlling for costs, and find similar levels of asymmetric pass-through: in each case, prices respond approximately twice as much to the VAT increase as to the VAT decrease (see table E7).

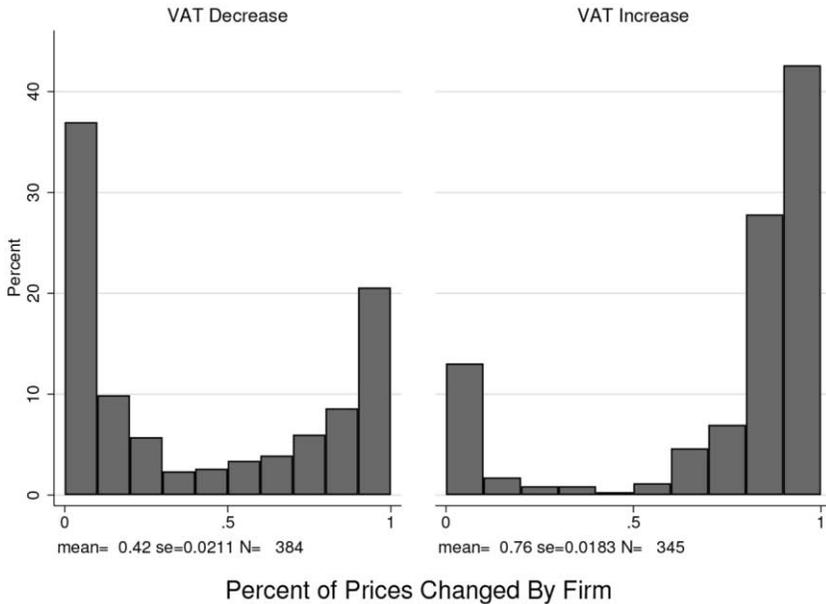


FIG. 2.—Proportion of prices changed by hairdresser. This figure plots the distribution of the within-hairdresser ratio of services for which prices were changed over total services offered following the VAT cut and hike.

difference of prices 1 month before and 1 month after the VAT reform: $\rho_i = \log(p_{\text{after}}) - \log(p_{\text{before}})$. Figures 4A and 4B plot the distribution of ρ_i for the VAT decrease and increase, respectively, and include all nine types of services.¹² The pass-through distribution for the VAT increase is bimodal: 27% of prices do not respond to the VAT increase, while 48% of prices increase by 80% to 120% of the amount of the VAT increase. The distribution of pass-through for the VAT decrease is unimodal: 61% of prices do not change in response to the VAT cut, while the rest decrease but without concentrating tightly at full pass-through levels (12% are located within 20% of full pass-through).

The asymmetry in pass-through distributions is not driven by specific services: we systematically observe a bimodal distribution following the VAT increase and a unimodal distribution following the VAT decrease for each of the nine services offered by hairdressers.¹³ Some of the observed heterogeneity could instead be explained by firm heterogeneity. In figure 2, we count the number of prices that are changed by firm, divide it by the number of services offered by each firm, and then plot the

¹² Figure E11 plots a version of fig. 4 that controls for inflation. The distributions are very similar but otherwise shifted to the left.

¹³ See figs. E12–E15.

distribution of the resulting ratio. The distributions are bimodal, which suggests the presence of two types of firms: those that tend to change all prices and those that keep all prices fixed. This finding is consistent with the argument made by Kopczuk and Slemrod (2006) and Slemrod and Gillitzer (2013), who point to the importance of accounting for firm-level heterogeneity when modeling tax behavior.

2. Pass-Through Distribution Dynamics in the Medium Run

To explore the dynamics of the pass-through distributions, we use a different price data set collected by Statistics Finland that allows us to observe prices for longer horizons. This data set has three main drawbacks. First, prices are unlikely to be randomly collected.¹⁴ Second, the data set does not allow us to observe prices immediately after the reform.¹⁵ Third, this data set stops in 2010, thus not allowing us to use it to analyze the VAT increase.

Figure 3 plots the distribution of $\rho_t = \log(p_t) - \log(p_0)$, where p_0 is price 1 month before the VAT change and p_t is price measured $t = \{6, 12, 18, 24\}$ months after the VAT decrease. These figures show that most of the price adjustments occur within a 6 month window, after which the excess mass of inert prices remains constant. The distributions look qualitatively similar to the short-run distributions from figure 4, with a mass at zero and the remaining price changes being negative without concentrating at full pass-through. However, the size of the spike at zero, while constant over time, is different from the one we observe in the short run in figure 4. This is likely due to the fact that the Statistics Finland data set does not randomize the collection of prices, whereas the one used in figure 4 does.

C. *Asymmetric Response of Profits and Markups*

Using administrative corporate tax data on the full population of hairdressers and beauty salons, we next investigate the response of profits and markups to VAT changes. In this data we observe turnover, profits, and variable and fixed costs, among other variables. As a proxy for markups, we use turnover minus variable cost divided by variable cost. This proxy is accurate as long as marginal costs are constant.

¹⁴ Statistics Finland oversamples larger firms and firms with prices that are easy to collect, such as firms with online prices and firms in the Helsinki area.

¹⁵ The price data collection is such that not all observations are updated immediately, and it can take up to 6 months for a given price to be updated.

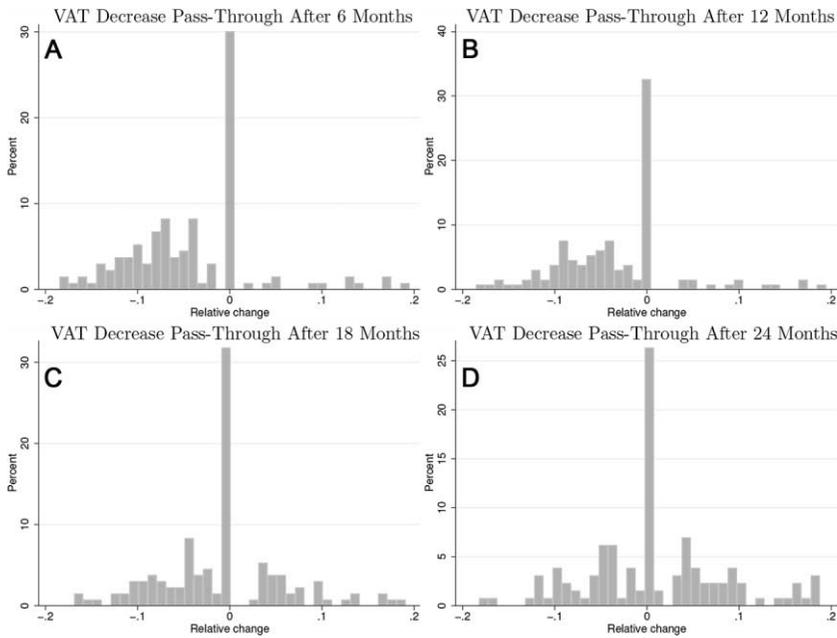


FIG. 3.—Medium-term pass-through distribution of the Finnish hairdressing VAT cut. These figures plot the distribution of the pass-through of the Finnish hairdressing VAT cut over four different horizons: 6, 12, 18, and 24 months. Pass-through is calculated as $\log(p_t) - \log(p_0)$, where p_0 is the price 1 month before the reform and p_t is the price $t = 6, 12, 18,$ and 24 months after the reform.

Figure 5A plots the coefficients from a regression of log profits on year dummies from 2000 to 2014 for hairdressers and beauty salons.¹⁶ The graph shows that profits respond asymmetrically to the VAT changes: the VAT decrease results in an increase in profits of 0.2 log points, while the VAT increase leads to a profit decrease of 0.1 log points. Figure 5B shows a similar graph for markups, which increase by twice as much following the VAT decrease as they decrease following the VAT increase. We observe no evidence of convergence of profits or markups toward symmetry 3 years after the VAT reverts to its original level.¹⁷ In contrast, as shown in figures 5C and 5D, we observe no significant changes in variable and fixed costs following the VAT changes, thus suggesting that quantities are not affected by the reform. These observations are consistent with firms using VAT decreases to

¹⁶ We exclude firms with less than €10,000 in turnover or €1,000 in profits to exclude small firms that are exempt from remitting VAT.

¹⁷ Note that we detect no evidence of increased entry (or exit) in the hairdressing sector following the VAT changes, as shown in figs. E16 and E17, which plot the number of firm entries and exits over time, respectively. Similarly, investment does not respond to the reforms, as shown in fig. E18.

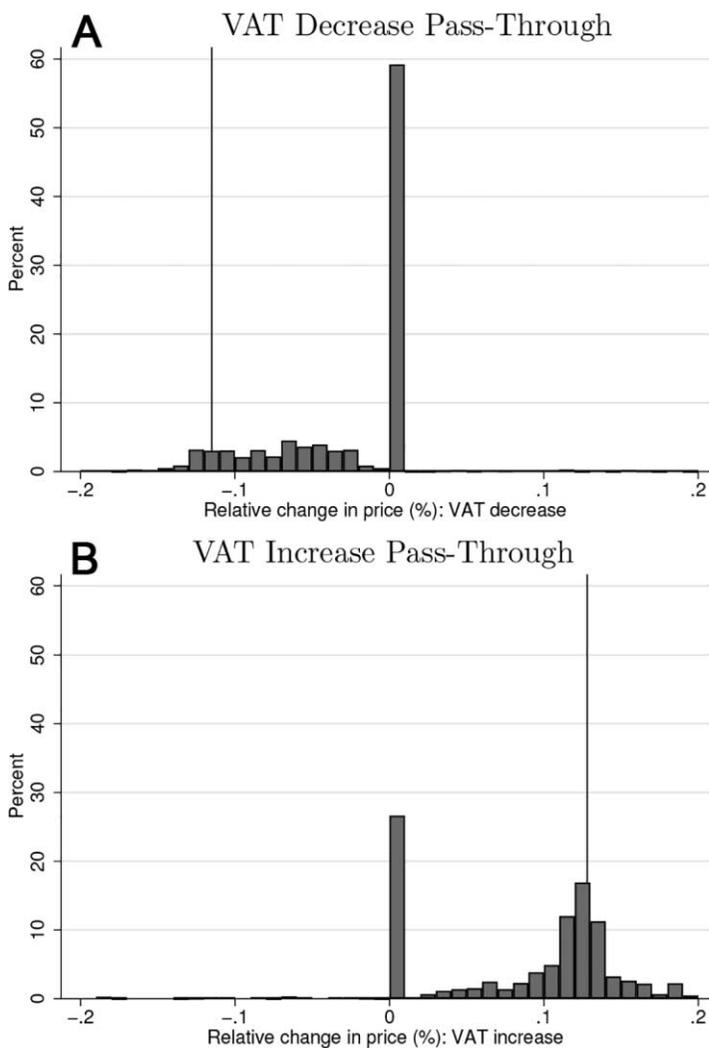


FIG. 4.—Distributional asymmetry. These figures plot the pass-through distributions of the VAT decrease (A) and the VAT increase (B) for Finnish hairdressing services. The vertical line represents 100% pass-through.

increase profits, while passing through VAT increases to prices to minimize their negative impacts on profits.

D. *Heterogeneous Firm Response*

The main dimension of heterogeneity we uncover is that firms with low profit margins at the time of the VAT change tend to pass through more

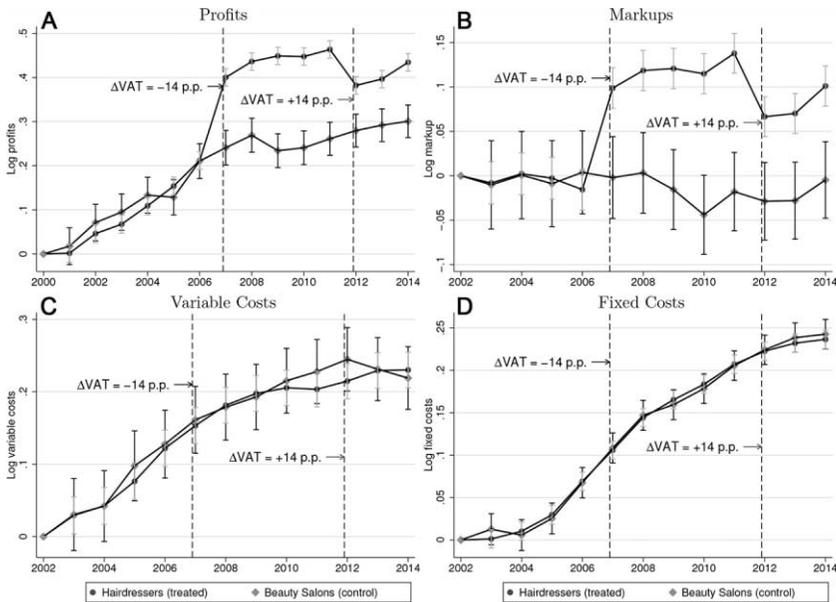


FIG. 5.—Plots of the coefficients from a regression of log profits (A), log markups (B), log variable costs (C), and log fixed costs (D) on year dummies for Finnish hairdressers and beauty salons.

of the VAT increase than of the decrease, whereas firms with high profit margins are more likely to behave symmetrically. Note that since profit margins are not economic primitives, this difference could be masking underlying firm differences that are not captured by our data. We define profit margins as turnover minus operating costs divided by turnover. To mitigate concerns of mean reversion, we calculate a 3 year average profit margin prior to the first VAT change (from 2004 to 2006) and break down our sample of hairdressers into quintile groups from the lowest profit margin to the highest.¹⁸ Figure 6A plots the change in markup in 2007 and 2012 for each quintile of profit margins, and shows that hairdressers in the lowest quintile take advantage of the VAT cut to increase their markups whereas firms in higher quintiles tend to behave symmetrically. To further mitigate concerns of mean reversion, we carry out two placebo tests: (1) figure 6B plots the response of the same quintiles in 2006 and 2011 and (2) figure 6C plots the response of beauty salons in 2007 and 2012. We find that changes in markups are significantly more homogeneous across quintiles in the placebo tests.

¹⁸ Figure E19 performs the same test using a different definition of the quintile margins by using the 2004 to 2006 data for the 2007 reform and the 2009 to 2011 data for the 2012 reform.

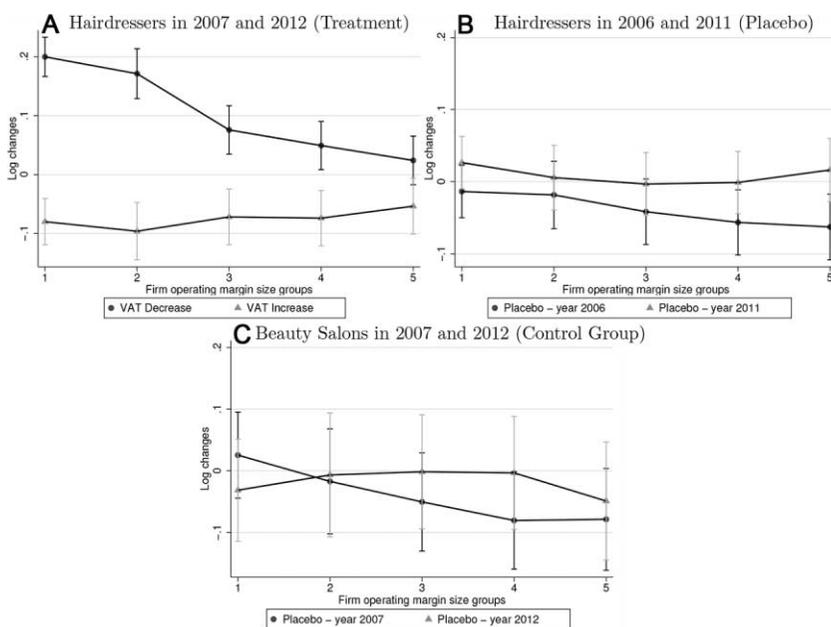


FIG. 6.—Changes in markups by quintile of profit margins for the Finnish VAT reforms. These graphs are generated by breaking down the sample of firms into quintiles of profit margins (turnover minus deductible costs divided by turnover) using data from 2004 to 2006, with quintile 1 being firms with the smallest profit margins. For each quintile, we plot changes in markup following changes in the VAT rate. *A*, The 14 percentage point VAT increase and decrease for Finnish hairdressers. *B*, Finnish hairdressers in 2006 and 2011. *C*, Finnish beauty salons (which we use as a control group for hairdressers) in 2007 and 2012.

E. What Do We Learn from the Finnish Evidence?

The Finnish VAT experiment has several advantages: (1) the timing and size of the VAT changes and the commodities affected by the reforms are not chosen by the Finnish government but imposed by the European Commission, (2) the VAT increase and decrease are of the same magnitude and are large (14 percentage points), and (3) we have rich firm level and price data that allow us to precisely assess the effects of these changes.

However, this experiment suffers from two main shortcomings. First, it is difficult to draw statistical inferences on the magnitude of the asymmetry if one only observes two VAT changes. Second, while the VAT increase and decrease affect the same commodities in the same country and are of the same magnitude, they do not occur at the same time. More importantly, the VAT decrease occurs prior to the Great Recession and the VAT increase occurs after. Reassuringly, we do not observe any evidence that economic conditions differentially affect our treatment and control groups: there is no indication of sharp changes in prices (fig. 1), profits (fig. 5*A*), markups (fig. 5*B*), or costs (figs. 5*C* and 5*D*) in either case except at the

time of the VAT changes. Nonetheless, we cannot rule out the possibility that differences in economic conditions between 2007 and 2012 might be affecting pass-through rates in a way that could bias them toward asymmetry.

For this reason, we turn to an analysis of all VAT changes occurring in the European Union from 1996 to 2015. Using these data, we can draw statistical inferences on the magnitude of the asymmetry, since the sample of reforms is substantially larger and covers many countries, periods, and commodities. Moreover, we can, to some extent, repeat the Finnish experiment using a matching approach that we explain in detail below.

V. European VAT Changes

A. Graphical Evidence

First, we provide some graphical and nonparametric evidence by examining the unconditional mean of the price index using the full sample of EU VAT changes in the 3 months before and after a VAT reform. In each case, we normalize the series to 100 in the month preceding the reform.

Figure 7A plots the unweighted average price of all commodities separately for VAT increases and decreases. The figure shows that prices increase discontinuously in the month following a VAT increase but decrease less in response to a VAT decrease. Importantly, this asymmetry is not driven by any subset of commodities. Indeed, when we plot disaggregated versions of figure 7A by two-digit COICOP groups (in figs. E20–E22; figs. E11–E23 are available online), we find that *all* commodities exhibit this asymmetric pass-through, with the exception of communication (COICOP group number 8) and furnishings, household equipment, and routine household maintenance (COICOP group number 5).¹⁹

B. Fixed-Effects Estimates

To estimate the pass-through to prices of VAT increases and decreases, we first follow the approach of Evans, Ringel, and Stech (1999), who estimate the pass-through of cigarette taxes using different tax changes across US states over time. We run the following fixed-effects regression:

$$\begin{aligned} \Delta \log(p_{ict}) &= \beta_0 \Delta \log(1 + \tau_{ict}) \\ &+ \sum_{k=-10, k \neq 0}^{k=10} \beta_k \Delta \log(1 + \tau_{ic(t+k)}) + \Delta \lambda_t + \gamma \Delta X_{ct} + \Delta \epsilon_{ict}, \end{aligned} \quad (2)$$

¹⁹ Possibly because of sample size, trends do not appear to be parallel for two COICOP categories: (1) alcoholic beverages, tobacco, and narcotics and (2) transport. While pass-through is asymmetric in both of these cases, the violation of the parallel assumptions suggest that we should be cautious in interpreting these two figures.

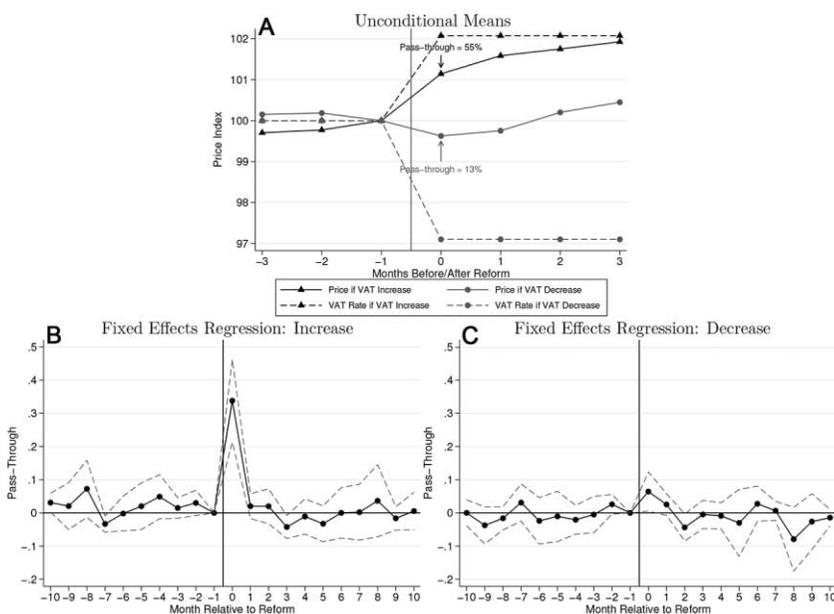


FIG. 7.—Asymmetric response of prices to VAT changes. *A*, Response of the unconditional means of prices to VAT increases and decreases. Lower panels plot the coefficients from the fixed-effects regression (2) for VAT increases (*B*) and VAT decreases (*C*) on the full sample of reforms and include 10 month leads and lags.

where i denotes the commodity, c the country, and t the month in which the price is observed, λ_t time fixed effects, p_{ict} the price, τ_{ict} the VAT rate, and ε_{ict} the error term. We control for a given country's nominal interest rate, GDP per capita, and unemployment rate with X_{ct} . For each of $x_t \in \{\log(p_{ict}), \log(1 + \tau_{ict}), \lambda_t, X_{ct}, \varepsilon_{ict}\}$, Δx_t is equal to $x_t - x_{t-1}$.

In equation (2), β_0 identifies the pass-through of a VAT change in the month when the change occurs: for example, if $\beta_0 = 0$, then the price does not respond to a VAT change, and if $\beta_0 = 1$, the price responds one-to-one to a VAT change. The second term of the equation estimates any forward- or backward-looking responses of prices to changes in VAT rates; β_{-5} , for example, identifies the response of prices at time t to VAT changes that will occur at time $t + 5$.

The fixed-effects regression generalizes an event-study design with multiple periods, commodities, and countries, and its main identification assumption is the same as that for difference-in-differences regressions: absent the tax change, there would have been no change in the prices of the treated relative to the untreated commodities. Figure 7A shows a sharp change in prices at the time of the reform, with no pretrends and no evidence of anticipatory behavior, which lends support to this identification

assumption. Identification relies on within-country-specific commodity variation in VAT rates over time.

In order to estimate the pass-through rate of VAT increases and decreases, we estimate equation (2) separately on two samples. The first sample includes all observations treated by a VAT increase along with all untreated observations and is used to estimate the pass-through rate of VAT increases. Similarly, the second sample includes all observations treated by a VAT decrease along with all untreated observations and is used to estimate the pass-through rate of VAT decreases.²⁰ Therefore, commodity-country pairs that never undergo a VAT change are included in both samples.

The results of these fixed-effects regressions are reported in table 1.²¹ Columns (1) and (2) of table 1 correspond to VAT increases and decreases, respectively. The first row of each regression (labeled β_0) shows the pass-through of the VAT change to prices 1 month after the reform. In particular, β_{+i} shows the response of prices to VAT changes i months after the reform, while β_{-i} shows the response of prices i months before the reform. Figures 7B and 7C plot the coefficients from regressions for VAT increases and decreases, respectively, and show that the pass-through to prices of VAT increases is equal to 34% while that of VAT decreases is equal to 6% one month after the reform; both are statistically significant. Importantly, there are no significant price responses in any months within a 10 month window around the VAT increases and decreases. We perform several robustness checks, including running specification (2) without controls but with time fixed effects (table E8; tables E5–E23 are available online) separately on reforms that are classified as temporary and permanent (tables E9 and E10), in periods of above- and below-median growth (tables E11 and E12), on reforms that are not concurrent with other tax changes (table E13), with country-commodity specific inflation controls (table E14), and including more leads and lags (figs. E23A and E23B). We find similar levels of asymmetry in each case. In addition, we examine the relationship between the timing of VAT changes with underlying economic conditions and find that the two are not correlated. These correlations are discussed in detail in appendix B.

²⁰ Untreated observations refer to commodity-country pairs that never undergo a VAT change as well as commodity-country pairs that are treated by a VAT change at a different time than the treatment period.

²¹ In a recent working paper, Abraham and Sun (2018) show that event study specifications can be problematic under certain conditions. In our case, these concerns would apply if there were heterogeneous effects across groups of VAT changes occurring at the same time. Given that our evidence is robust across a wide range of specifications as well as in several subsamples and also when plotting unconditional means in fig. 7A, this concern is unlikely to apply to our setting.

TABLE 1
PASS-THROUGH ESTIMATES USING FIXED-EFFECTS
REGRESSION (FULL SAMPLE)

| | $\Delta \log \text{PRICE}$ | |
|--------------------|----------------------------|-----------------|
| | Increase (1) | Decrease (2) |
| β_0 | .34 (.063) | .064 (.030) |
| β_{+1} | .020 (.019) | .025 (.017) |
| β_{-2} | .030 (.019) | .026 (.015) |
| β_{+2} | .020 (.027) | -.044 (.021) |
| β_{-3} | .015 (.016) | -.005 (.028) |
| β_{+3} | -.043 (.017) | -.004 (.022) |
| β_{-4} | .049 (.034) | -.020 (.022) |
| β_{+4} | -.011 (.027) | -.008 (.020) |
| Time fixed effects | Yes | Yes |
| R^2 | .014 | .014 |
| Observations | 386,557 | 342,792 |

NOTE.—The coefficients reported in this table show the pass-through of VAT increases and decreases to prices, estimated using specification (2) on the full sample of reforms. Column 1 shows the estimates for VAT increases, and col. 2 those for VAT decreases. Standard errors are clustered by month and are in parentheses. β_0 measures the pass-through of the VAT change at the time of the reform, and β_i measures price changes i months away from the reform.

C. Matching Estimates

In this section, we supplement our fixed-effects estimation with a matching approach that allows us to better mirror the setting of the Finnish experiment and to more directly interface with the theoretical predictions derived in section II. This approach also allows us to evaluate and alleviate concerns that in the preceding analysis we are never able to observe the effect of VAT increases and decreases at the same time, for the same commodity, and in the same country.

Our matching procedure follows Imbens (2015), and in particular that paper's approach to analyzing the Imbens, Rubin, and Sacerdote (2001) lottery data (in sec. 6.1). First, and motivated by our theoretical framework from section II, we choose a set of key characteristics used to estimate propensity scores. These characteristics include time, measures of economic activity, the magnitude of the VAT change, and the type of commodity

TABLE 2
PASS-THROUGH ESTIMATES: MATCHED SAMPLE

| | WITH TRIMMING | | WITHOUT TRIMMING | |
|--------------------|-----------------------------|------------------|-----------------------------|------------------|
| | $\Delta \log \text{ Price}$ | | $\Delta \log \text{ Price}$ | |
| | Increase (1) | Decrease (2) | Increase (3) | Decrease (4) |
| β_0 | .29 (.10) | .042 (.047) | .33 (.066) | .056 (.037) |
| β_{+1} | .021 (.019) | .025 (.017) | .020 (.019) | .025 (.017) |
| β_{-2} | .031 (.019) | .026 (.015) | .030 (.019) | .026 (.015) |
| β_{+2} | .021 (.027) | -.044 (.021) | .020 (.027) | -.044 (.021) |
| β_{-3} | .016 (.016) | -.0050 (.028) | .015 (.016) | -.0050 (.028) |
| β_{+3} | -.042 (.018) | -.0047 (.022) | -.042 (.018) | -.0046 (.022) |
| β_{-4} | .051 (.034) | -.020 (.022) | .049 (.034) | -.020 (.022) |
| β_{+4} | -.0092 (.027) | -.0089 (.020) | -.0096 (.027) | -.0089 (.020) |
| Time fixed effects | Yes | Yes | Yes | Yes |
| R^2 | .013 | .014 | .014 | .014 |
| Observations | 384,559 | 342,453 | 386,080 | 342,477 |

NOTE.—This table reports the pass-through estimates of VAT increases and decreases to prices estimated using specification (2) on matched reforms. Columns 1 and 2 implement the matching estimator on a trimmed sample using the approach outlined in Imbens (2015). Propensity scores are estimated using logit. Columns 3 and 4 report the estimates for the matched sample without trimming. VAT increases and decreases are matched on time, type of commodity (COICOP digit 2), size of VAT change, and GDP growth. In addition, we match the preincrease VAT rate to the postdecrease VAT rate and the postincrease VAT rate to the predecrease VAT rate. Columns 1 and 3 show the estimates for VAT increases, and cols. 2 and 4 those for VAT decreases. Standard errors are clustered by month and are in parentheses. β_0 measures the pass-through of the VAT change at the time of the reform, and β_i measures price changes i months away from the reform.

(two-digit COICOP code). Most importantly, guided by results from section II, we also match preincrease VAT rates to postdecrease VAT rates and predecrease VAT rates to postincrease VAT rates. Second, we use the estimated propensity scores, based on these key characteristics, to trim the sample as suggested by Crump et al. (2009). Finally, in this trimmed sample, we match the VAT increases and decreases, using the matching estimator outlined in section 5.4 of Imbens (2015).²² We then estimate equation (2) on the matched sample of VAT changes.

Estimates for pass-through from this matched sample, summarized in table 2, show that our finding of asymmetry is robust to matching. Indeed,

²² We cluster our standard errors as we do in our main specification. Bootstrapping yields similar standard errors, with all VAT increase pass-through estimates significant at least at the 1% level. However, we do not report these here, following cautions in Abadie and Imbens (2008) against using bootstrap with matching estimators.

these estimates reveal asymmetries that are similar to those estimated using the full (unmatched) sample of reforms. In appendix A, we also formally test the ability of our matching procedure to increase overlap in covariates. We follow Imbens (2015), who uses normalized differences in order to assess overlap. Our matching procedure substantially improves overlap and reduces all normalized differences. While there is no formal threshold for normalized differences and while our matching procedure does not reduce all normalized differences to zero, they are well in line with those estimated by Imbens (2015) using the Imbens, Rubin, and Sacerdote (2001) lottery data.

Because our matched sample mirrors the Finnish experiments, our pass-through estimates can be directly compared to the predictions of standard incidence theory discussed in section II. We conclude that our finding of asymmetric pass-through in the matched sample is difficult to reconcile with static incidence theory.

We provide more details on how the matching procedure is implemented in appendix A, and include formal tests of pre- and postmatching overlap in observables and further robustness checks. Note that a limitation of our approach is that we cannot match on prereform price levels because our data do not allow us to condition on price levels in a consistent way across all EU countries and commodities. This limitation implies that we are implicitly assuming that, conditional on matching on the key characteristics outlined above, any remaining variation in price levels has no meaningful effect over the asymmetry. For this reason, a stronger test of the asymmetry would also match on price levels.

D. Long-Term Persistence of the Asymmetry

In subsection IV.A we showed that asymmetry in the pass-through of VATs persisted for several years after the Finnish hairdressing reforms were enacted. In this section, we provide evidence suggesting that this persistence is not a peculiarity of Finnish hairdressers: we observe it in other sectors and countries as well. In order to analyze longer-run horizons, it is important to focus attention on relatively large VAT changes and sectors in which prices are relatively stable. Absent these conditions, the effect of VAT changes on prices might be masked by natural variation in prices and inflation that may be quantitatively significant in the long run.

In order to study this we start by considering two case studies that involve large VAT changes. First, the European Commission approved an application to reclassify sit-down restaurants from the standard VAT rate to the reduced VAT rate.²³ Both France and Finland took advantage of this

²³ Following a campaign promise by then French president Jacques Chirac, France applied for an authorization to reclassify sit-down restaurants from the standard to the reduced VAT rate in 2002. The application was approved for all member states in January 2009.

new law, leading to a 14 percentage point VAT cut for French sit-down restaurants and a 9 percentage point cut for Finnish ones, in 2009. While the VAT rate did not revert to its original level as in the Finnish hairdressing experiment, we can exploit smaller increases in the reduced VAT rate: 1.5 and 3 percentage point increases in France, in 2012 and 2014, and a 1 percentage point increase in Finland, in 2013. Figures 8A and 8B show that the asymmetric pass-through persisted over several years both in Finland and in France. This evidence is to be interpreted cautiously, in part because the VAT changes are not of the same size and it is conceivable that the asymmetric pass-through is in part due to that.

Second, and more comprehensively, Hungary cut its standard VAT rate from 25% to 20% in January 2006 and increased it from 20% to 25% in July 2009. These changes were enacted as part of a set of campaign promises preceding the 2006 parliamentary elections. Figure 8C shows the response of commodities that were subject to the standard rate in Hungary

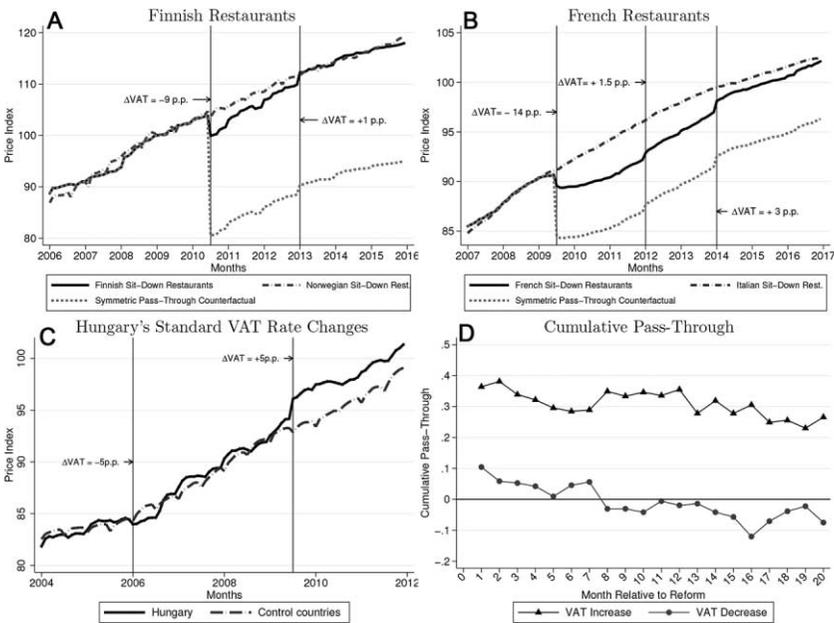


FIG. 8.—Long-term persistence of the asymmetry. *A*, Response of Finnish sit-down restaurants to a 9 percentage point VAT decrease and a 1 percentage point VAT increase compared to a control group of Norwegian sit-down restaurants. *B*, Response of French sit-down restaurants to a 14 percentage point VAT decrease and 1.5 and 3 percentage point VAT increases relative to a control group of Italian restaurants. We also include a hypothetical counterfactual that uses the VAT increase pass-through for VAT decreases. *C*, Response of all commodities subject to the standard VAT rate in Hungary (excluding diesel and gasoline) to a 5 percentage point VAT decrease and a 5 percentage point VAT increase relative to a control group consisting of neighboring countries. *D*, Cumulative pass-through rates of all VAT increases and decreases in our sample over 20 months.

compared to a set of control countries.²⁴ We find not only strong evidence of asymmetry but also that the asymmetry persisted over several years after the VAT rate was returned to 25%. Because the standard VAT rate applies to a wide range of commodities, this finding suggests that the long-term persistence of asymmetry may arise in more than a small set of specific sectors.

Finally, we can assess the long-term persistence of the asymmetry in the full sample of reforms. To do so, we first reestimate equation (2) with 20 leads and 20 lags (we plot the resulting coefficients in figs. E23A and E23B). We find no substantial lagged responses, even 20 months after the changes take place. Second, we add up all of the lag coefficients and plot the cumulative pass-through rate for VAT increases and decreases over a 20 month period following the VAT changes in figure 8D. While both cumulative pass-through rates appear to decrease over time, the magnitude of the asymmetry—which is given by the difference between the two cumulative pass-through rates—remains relatively constant and shows no signs of convergence. However, in spite of the two time series being relatively stable, the precision of the cumulative pass-through estimates mechanically decreases over time, making any long-run statements past 18 months only suggestive.²⁵

VI. Conclusion

In this paper, we show that prices respond asymmetrically to VAT changes. First, prices respond more to VAT increases than to VAT decreases. Second, this asymmetric response results in an asymmetric pass-through of VAT changes to profits and markups. Third, the asymmetry persists over the long run.

These patterns are inconsistent with the predictions of standard, static incidence models, and our findings therefore call into question the empirical relevance of these models for an important and widely employed tax. It is important for future research to examine how general these findings are and, in particular, to what degree the types of asymmetries we document arise with other taxes. More broadly, our findings suggest that it may be important to develop and empirically assess alternative models of incidence that can account for these types of asymmetries. One potentially promising avenue is to move beyond the static setting of classical

²⁴ We included every commodity subject to the standard VAT rate, with the exception of diesel and gasoline because of strong volatility. Details of the list of commodities and control group countries can be found in app. D (apps. C–E are available online).

²⁵ We cannot reject that the cumulative VAT increase and decrease pass-through rates are equal after 18 months. We report the level difference between the two cumulative pass-through rates and the corresponding significance levels in table E15.

models, and evaluate to what degree *dynamic* models of tax incidence can generate asymmetries like the ones we document.

Appendix A

Matching VAT Increases and Decrease

In this appendix, we first provide details of the matching approach we use in section V.C and then discuss and provide the results of several alternative matching algorithms and specifications.

A1. Matching Estimator

Our matching approach follows Imbens (2015). We follow the application of its approach to the Imbens, Rubin, and Sacerdote (2001) lottery data, outlined in section 6.1 of Imbens (2015).²⁶ The matching procedure outlined in Imbens (2015) follows three steps. First, it estimates propensity scores in the original sample using a logistic regression. Second, it trims the original sample by dropping extreme observations, using the threshold derived in Crump et al. (2009; referred to as CHIM), by removing any observations with extreme propensity scores. Third, it implements a matching algorithm—which is described in detail in section 5.4 of Imbens (2015)—on the resulting sample. We implement these three steps on our sample of VAT increases and decreases and then reestimate specification (2) on the resulting subsample of VAT reforms.

Note that we report the resulting estimates for both the trimmed and matched sample (which corresponds to the Imbens [2015] approach) as well as the matched sample without trimming (which skips the second step in Imbens [2015]). We do this because the trimmed specification has an even more stringent definition of poor matches compared to the matched-only specification. Indeed, it drops observations with extreme propensity scores that would otherwise be kept in the matched-only specification. Assessing the difference in estimates between the no-matching specification (table 1), the matched-only specification (table 2, cols. 3 and 4), and the matched and trimmed specification (table 2, cols. 1 and 2) should give us a sense of how likely the estimated asymmetry is due to VAT increases and decreases being different. If pass-through estimates for VAT increases and decreases become closer as the sample of VAT increases and decreases is made more comparable by matching and trimming, one should be seriously concerned as to whether the estimated asymmetry is mechanically driven by the VAT changes being inherently different. If instead, the estimated asymmetry is similar, concerns that differences in the samples of VAT increases and decreases are driving the asymmetry would be seriously mitigated.

Table 2 reports the results of this matching estimator. Overall, our finding of asymmetry is robust both in the matched sample (cols. 3 and 4) and in the matched and trimmed sample (cols. 1 and 2). The more stringent specification,

²⁶ We rely on the first example, because it is the closest to our setting as the second and third examples rely on experimental data.

which is the trimmed and matched sample, generates levels of asymmetry that are similar to the matched-only sample or the full sample, making a strong case against the concern that the differences between the VAT increase and decrease samples are driving the asymmetry. Nevertheless, we conduct several robustness checks below.

A2. *Assessing the Performance of the Matching Estimator*

Imbens (2015) suggests using normalized differences rather than using t -statistics or visually comparing distributions, in order to assess overlap between the matched samples. Following the notation of Imbens (2015), the formula for normalized differences is given by

$$\Delta_{X,k} = \frac{\overline{X_{inc,k}} - \overline{X_{dec,k}}}{\sqrt{(\overline{S_{X,inc,k}^2} + \overline{S_{X,dec,k}^2})/2}}, \quad (A1)$$

where $\overline{X_{dec,k}} = (1/N_{dec})\sum_{i \in D} X_{i,k}$, $\overline{X_{inc,k}} = (1/N_{inc})\sum_{i \in I} X_{i,k}$, $\overline{S_{X,dec,k}^2} = [1/(N_{dec} - 1)]\sum_{i \in D} (X_{i,k} - \overline{X_{dec,k}})^2$, and $\overline{S_{X,inc,k}^2} = [1/(N_{inc} - 1)]\sum_{i \in I} (X_{i,k} - \overline{X_{inc,k}})^2$, where i is a given VAT change, $X_{i,k}$ is an element of the covariate vector X_i , N_{inc} is the number of VAT increases, N_{dec} is the number of VAT decreases, I is the set of VAT increases, and D is the set of VAT decreases.

Imbens (2015) argues that normalized differences are better suited than t -statistics for assessing overlap between two samples because t -statistics can be large when the sample size is large, rejecting that two sample means are equal even when they are not substantively different. As written in Imbens (2015), “The normalized differences provide a scale and sample size free way of assessing overlap.”

Larger normalized differences indicate larger average differences in covariates in the two groups that are being compared. Importantly, there is no specific “threshold” above which normalized differences are considered to be too large. Instead, we benchmark our normalized differences against those estimated by Imbens (2015) for the lottery data. In addition to assessing the level of normalized differences, another useful test is to compare the pre- and post-trimming normalized differences. A well-performing matching procedure should reduce any large normalized differences.

Table B2 reports the normalized differences for the main covariates that we consider in our main specification: time, magnitude of VAT change, economic conditions, commodity type, preincrease VAT rate matched to postdecrease VAT rate, and postincrease VAT rate matched to predecrease VAT rate. Note that rather than reporting normalized differences for the 235 different months, and since the main concern with timing is that reforms might happen in times of different economic conditions, we report instead normalized differences for a variable that is equal to 1 if the month is a month during which the economy is in a recession and 0 otherwise.

Two important points are worth emphasizing: (1) the normalized differences in the trimmed sample are small and, relatedly, (2) trimming substantially reduces any large normalized differences estimated in the original sample. On point 1, while there is no specific threshold for normalized differences, Imbens (2015) mentions 0.3, in absolute value, for example as a possible benchmark in section 6.1.1.

The normalized differences we estimate in the trimmed sample are all smaller than 0.3. Moreover, the normalized differences we estimate are smaller than the ones Imbens (2015) estimates using the lottery data. For example, two covariates in Imbens (2015) have normalized differences in excess of 0.3 in its trimmed sample (0.51 and -0.47), while the largest normalized difference in our trimmed sample is 0.27, in absolute value. On point 2, the variables that have large normalized differences in the full sample of VAT reforms are substantially affected by trimming: the normalized differences for economic conditions are reduced from 0.43 to 0.10, for predecrease/postincrease from -0.60 to 0.06, for preincrease/postdecrease from -0.33 to 0.07, for size of VAT change from -0.32 to -0.03 , and for recession months from 0.56 to 0.27. The remaining normalized differences for commodity types are small even before trimming.

Overall, trimming performs well, making the sample of VAT increases and decreases plausibly comparable. In spite of this substantial reduction in differences between the sample of VAT increases and decreases, we still estimate that pass-through is larger for VAT increases relative to VAT decreases, mitigating concerns that the asymmetry is an artifact of VAT increases and decreases being inherently different.

A3. *Further Robustness*

We implement two sets of further robustness checks and one regression correction that should account for any remaining imbalance and assess the plausibility of unconfoundedness.²⁷

First, we ensure that our results are robust to the variables we match on. Imbens (2015) automatizes the choice of matching variables by using an algorithm as well as lasso. We do not follow this approach for two reasons: (1) matching on the variables we choose performs very well at reducing normalized differences, and (2) the variables we choose are driven by specific concerns derived from theory (sec. II) and we believe these concerns provide a better justification than using automatic algorithms; Heckman, Ichimura, and Todd (1997), for example, insists on using economic theory to choose what variables need to be included when implementing a matching estimator. Nevertheless, we ensure that our results are not sensitive to variable choices and implement several additional specifications. First, we match on unemployment rate, in addition to annual growth, in table E17. Second, we exclude all matching variables and match only preincrease VAT rates to postdecrease VAT rates and postincrease VAT rates to predecrease VAT rates in table E18. While the estimated pass-through rates vary slightly, we systematically estimate similar levels of asymmetric pass-through, mitigating concerns that our main matching specification is not robust to the choice of matching variables.

Second, Imbens (2015) insists (in sec. 5.2) on testing whether the results are sensitive to estimating the propensity score using probit instead of logit. In table E19, we reestimate our main matching specification using probit and find similar levels of asymmetric pass-through.

²⁷ We also estimate eq. (2) where, in addition to the baseline terms, we include interactions of the treatment variable with all covariates and find similar levels of asymmetric pass-through, as shown in table E16.

Third, Imbens (2015) points out that “unless a particular estimator is robust to modest changes in the implementation, any results should be viewed with suspicion.” Therefore, in addition to the estimator used by Imbens (2015) in the lottery data, we implement three more matching algorithms with the same baseline matching characteristics. Table E20 shows the pass-through estimated using these alternative algorithms. We find very similar levels of asymmetric pass-through.

Fourth, Imbens (2015) proposes a way of assessing whether unconfoundedness is plausible, which is to estimate the same baseline specification on the trimmed sample but instead of using the actual outcome, using a lagged outcome. The test passes if the estimated coefficient is small and not statistically significant. We implement this test and estimate a coefficient on our lagged outcome of 0.0033 (0.014) for VAT increases and 0.036 (0.026) for VAT decreases. Given the small magnitude of the estimates and the fact that none are statistically significant, unconfoundedness seems plausible.

Fifth, Imbens (2015) proposes a method to correct for any differences in the two samples that remain after the samples are trimmed. The approach is explained in detail in section 5.4 of Imbens (2015). Table E21 shows the result of implementing this correction to our estimate. The magnitude of the asymmetry is very similar to our main matching specification, which is not surprising, given that the normalized differences we estimate in the trimmed sample are small.

Appendix B

The Timing of VAT Changes Is Not Correlated with Economic Conditions

This appendix implements an alternative to the matching approach (app. A) to address the concern that some of the EU VAT changes might be endogenous to the underlying economic conditions. First, we use economics and institutional knowledge to identify the variables we expect, *ex ante*, to be correlated with the timing of VAT changes. Second, we empirically test the correlation between the timing of VAT changes and the variables identified and find no evidence supporting it.

Ex ante, and based on our analysis of the underlying reasons for VAT changes and institutional details below, we can expect VAT changes to occur (1) for political reasons, such as electing more fiscally conservative governments; (2) for institutional reasons, mainly because of the VAT harmonization efforts led by the European Union; and (3) for economic reasons, such as using VAT changes to counteract changing economic conditions. We describe each of these below.

First, there are reasons to expect the timing of VAT changes to be correlated with political variables, such as the strength of the governing coalition and changes in the governing party. There is an empirical political economy literature that analyzes the underlying reasons for tax reforms and finds that political reasons are more likely to cause tax reforms than economic conditions. Castanheira, Nicodème, and Profeta (2012), for example, show that political variables (strength of the governing coalition and weakness of the opposition party) are more likely to predict tax reforms than economic conditions (GDP and unemployment). Moreover, Hallerberg and Scartascini (2017) show that electoral considerations are more

likely to drive VAT changes than economic considerations. Moreover, Foremny and Riedel (2014) show that changes in local business taxes in Germany are driven by the electoral cycle. In principle, political changes could lead to other changes, which in turn could affect the underlying economic conditions. For this reason, we perform several tests to ensure that economic conditions are unlikely to explain the asymmetry and discuss them in the third point below.

Second, VAT changes could also be due to institutional reasons. The European Commission adopted legislation in 2006 that significantly restricted the ability of member states to freely set their VAT rates. Council Directive 2006/112/EC explicitly mandated that member states should progressively start abiding by the following rules: (1) increase the standard VAT rate above 15% and the reduced VAT rate above 5%; (2) restrict the reduced VAT rate to a prespecified set of commodities, essentially preventing member states from artificially reducing VAT rates by reclassifying commodities from the standard to the reduced VAT rate; (3) any reduction of VAT rates below 15% (or reclassification from 15% to 5%) was to be approved by all 28 member states.²⁸ Given these restrictions, we can expect the following three patterns, after 2006: (1) fewer VAT decreases, (2) VAT decreases of smaller magnitude, and (3) more VAT increases aimed at bringing VAT rates above the 5% and 15% minima.

Finally, the timing of VAT changes could be correlated with economic conditions. Except for Council Directive 2006/112/EC, there are no other laws that restrict member states from using VATs to affect the economy. This could threaten our identification if VAT increases occur at times when economic conditions are particularly different from those of VAT decreases *and* prices respond differently to VAT shocks during those different times. Since this could threaten our identification strategy, we explicitly test for the correlation of economic conditions with the timing of VAT changes.

To do so, we estimate the correlation of the timing of increases and decreases with measures of economic conditions. To proxy for economic activity, we follow the National Bureau of Economic Research's Business Cycle Dating Committee, which, in the United States, is the organization that dates recessions and expansions. The main measures they consider are GDP and employment. The underlying reasoning is that GDP rises during periods of expansions, while unemployment falls, and conversely GDP falls during recessions while unemployment rises. Fuest, Peichl, and Siegloch (2018), for example, estimate the incidence of corporate taxes using changes in corporate tax rates over time, and use GDP and the unemployment rate to show that corporate tax changes are not driven by economic conditions. We use a similar approach, and test for the correlation of VAT changes with total tax revenue, GDP per capita, and the unemployment rate in the 12 months leading to a given VAT reform. We find no significant relationship between VAT reforms and these measures of economic activity. Formally, we run the following regression:

²⁸ The third rule is not explicitly laid out in Council Directive 2006/112/EC, but, procedurally, any exemptions to a given Council Directive require a vote by the 28 member states.

$$\begin{aligned} \text{reform}_{ict} = & \sum_{t=-12}^{t=-1} \alpha_t \log(\text{TR}_{ct}) + \sum_{t=-12}^{t=-1} \beta_t \log(\text{GDP}_{ct}) + \sum_{t=-12}^{t=-1} \gamma_t \log(\text{UR}_{ct}) \\ & + \lambda_t + \delta_c + \pi_i + \epsilon_{ict}, \end{aligned} \quad (\text{B1})$$

where reform_{ict} is equal to 1 if a VAT change occurs for commodity i in country c in month t and 0 otherwise; TR_{ct} is total tax revenue for country c in month t ; GDP_{ct} is the per capita GDP of country c in month t ; UR_{ct} is the unemployment rate of country c in month t ; λ_t are time (in months) fixed effects; δ_c are country fixed effects; π_i are commodity fixed effects; and ϵ_{ict} is the error term (clustered by month). We run this regression on the full sample, on a subsample excluding VAT decreases, and on a subsample excluding VAT increases. The outcome variable for the full sample is equal to 1 if there is a VAT change and zero otherwise; the outcome variable for the sample excluding VAT decreases is equal to 1 if there is a VAT increase and zero otherwise; and the outcome variable for the sample excluding VAT increases is equal to 1 if there is a VAT decrease and zero otherwise.

Table E22 shows that there is no relationship between the timing of VAT changes—whether increases or decreases—with the underlying economic conditions leading up to the reforms or with total tax revenue. Using sector-specific measures of economic conditions instead of GDP, such as turnover by sector, yields similar results, as shown in table E23. This further mitigates concerns that VAT changes are endogenous to economic conditions. In an additional robustness check, we also reestimate equation (2) separately for VAT changes that occur during periods of above- and below-median growth (tables E11 and E12) and find that pass-through rates are also asymmetric.

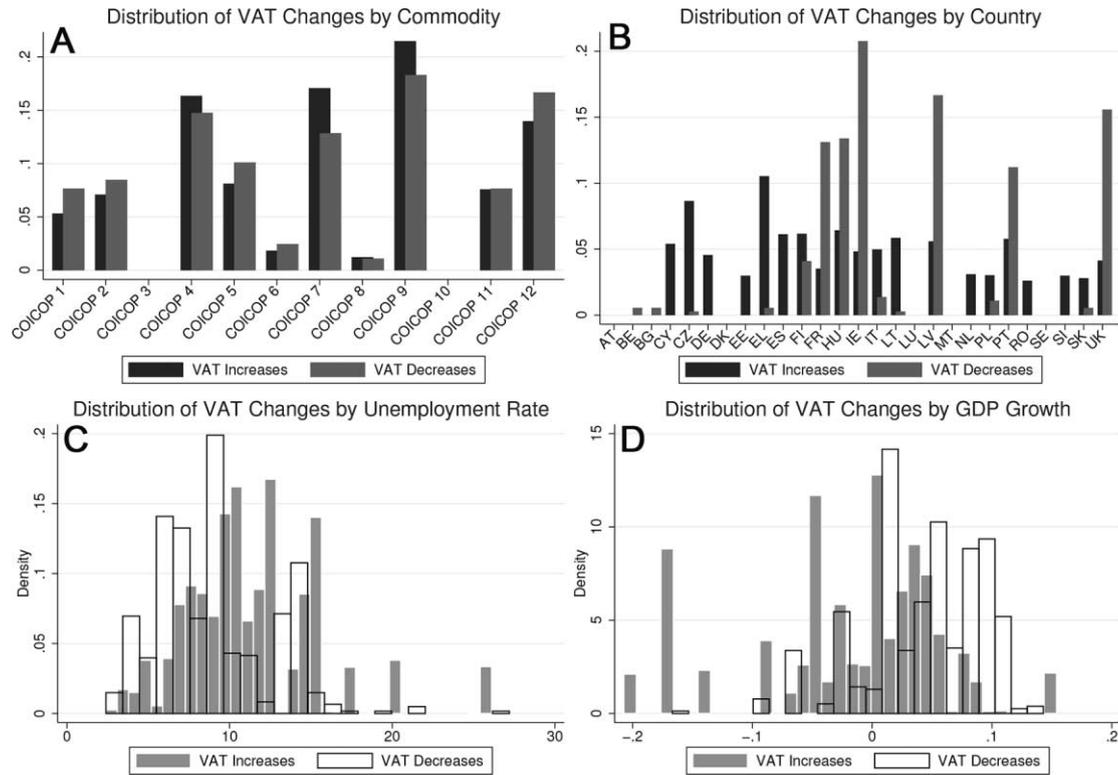


FIG. B1.—EU VAT changes: Summary statistics. These panels show the distribution of VAT increases and decreases by two-digit COICOP category (A), country (B), unemployment rate (C), and GDP growth rate (D). The description of the two-digit COICOP categories is provided in tables E5 and E6.

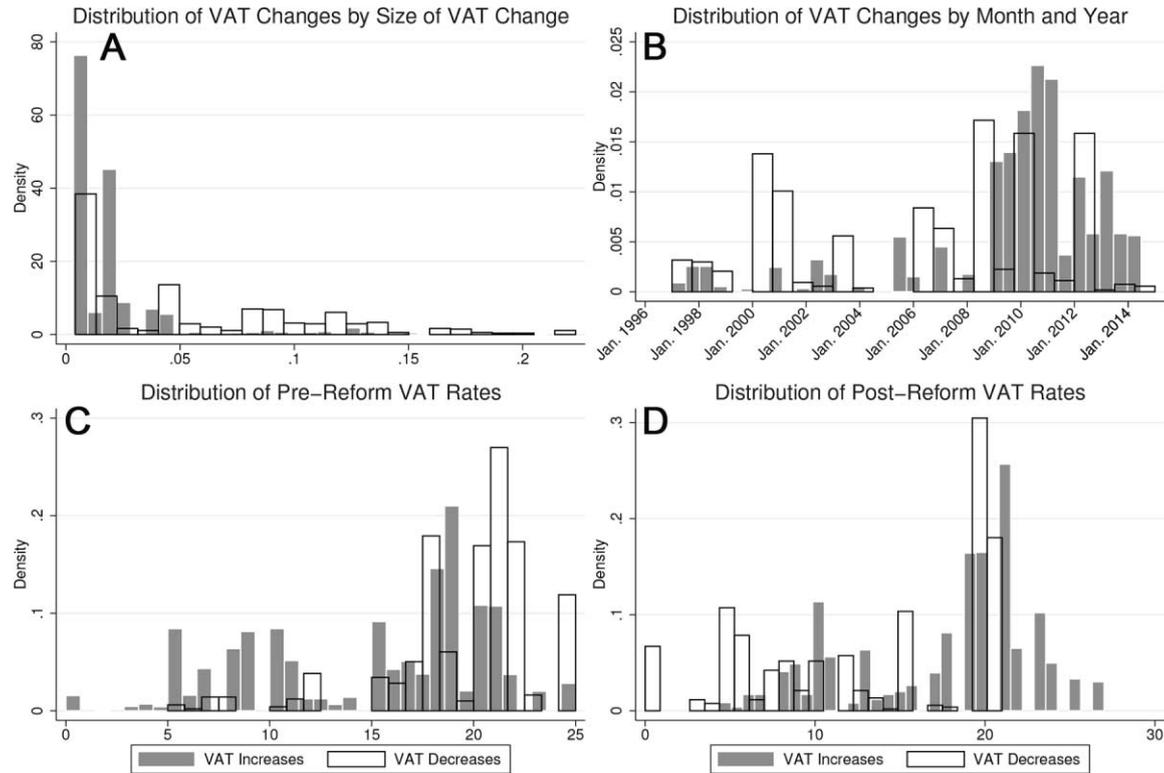


FIG. B2.—EU VAT changes: Summary statistics (continued). These panels show the distribution of VAT increases and decreases by size of the VAT changes (A), month and year (B), prereform VAT rate (C), and postreform VAT rate (D).

TABLE B1
SUMMARY STATISTICS FOR FINNISH HAIRDRESSERS AND BEAUTY SALONS

| | HAIRDRESSERS | | | | BEAUTY SALONS | | | |
|-------------------------|--------------|--------|--------------------|--------------|---------------|--------|--------------------|--------------|
| | Mean | Median | Standard Deviation | Observations | Mean | Median | Standard Deviation | Observations |
| Turnover | 40,190 | 25,924 | 231,039 | 157,082 | 35,643 | 18,504 | 143,747 | 45,368 |
| Profits | 13,787 | 11,330 | 15,193 | 155,837 | 9,610 | 5,048 | 19,365 | 44,332 |
| Costs | 26,699 | 13,285 | 213,093 | 162,634 | 26,865 | 11,415 | 126,093 | 47,347 |
| Total assets | 12,841 | 2,834 | 79,027 | 112,682 | 13,065 | 2,115 | 84,635 | 36,984 |
| Number of employees | .40 | 0 | 4.22 | 162,634 | .37 | 0 | 3.53 | 47,347 |
| Cost of employees | 1,129 | 0 | 20,138 | 145,729 | 766 | 0 | 10,709 | 43,649 |
| Sole proprietors | .91 | 1 | .29 | 162,634 | .89 | 1 | .31 | 47,347 |
| Partnerships | .05 | 0 | .21 | 162,634 | .03 | 0 | .18 | 47,347 |
| Corporations | .05 | 0 | .21 | 162,634 | .07 | 0 | .26 | 47,347 |
| Number of firms in 2006 | | | 12,301 | | | | 3,073 | |

NOTE.—This table reports annual summary statistics on the full population of Finnish hairdressers and beauty salons using corporate tax data.

TABLE B2
NORMALIZED DIFFERENCES BEFORE AND AFTER MATCHING

| | NORMALIZED DIFFERENCES | |
|--------------------------|------------------------|-----------------------|
| | Full Sample (1) | Trimmed Sample (2) |
| Preincrease-postdecrease | -.33 | .07 |
| Predecrease-postincrease | -.60 | .06 |
| Economic conditions | .43 | .10 |
| Size of VAT change | -.32 | -.03 |
| Recession months | .56 | .27 |
| COICOP 1 | -.09 | -.08 |
| COICOP 2 | -.05 | .00 |
| COICOP 3 | .04 | .05 |
| COICOP 4 | -.07 | -.02 |
| COICOP 5 | -.04 | .00 |
| COICOP 6 | .12 | -.05 |
| COICOP 7 | .01 | -.03 |
| COICOP 8 | .08 | .02 |
| COICOP 9 | .00 | .07 |
| COICOP 10 | -.07 | .01 |

NOTE.—This table reports the normalized differences for the variables we use to match VAT increases to VAT decreases using our main matching specification. Column 1 reports the normalized differences for the full sample of reforms, and col. 2 for the trimmed sample of reforms. Preincrease-postdecrease matches preincrease VAT rates to postdecrease VAT rates, and predecrease-postincrease matches predecrease VAT rates to postincrease VAT rates.

References

- Abadie, Alberto, and Guido W. Imbens. 2008. "On the Failure of the Bootstrap for Matching Estimators." *Econometrica* 76 (6): 1537–57.
- Abraham, Sarah, and Liyang Sun. 2018. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." Working paper, Dept. Econ., Massachusetts Inst. Tech.
- Atkinson, Anthony B. 2012. "The Mirrlees Review and the State of Public Economics." *J. Econ. Literature* 50 (3): 770–80.
- Benedek, Dora, Ruud A. De Mooij, Michael Keen, and Philippe Wingender. 2020. "Varieties of VAT Pass Through." *Internat. Tax and Public Finance* 27:890–930.
- Benzarti, Youssef, and Dorian Carloni. 2019. "Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France." *American Econ. J.: Econ. Policy* 11 (1): 38–63.
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert. 1997. "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Q.J.E.* 112 (1): 305–39.
- Carbonnier, Clément. 2008. "Différence des Ajustements de Prix à des Hausses ou Baisse des Taux de la TVA: Un Examen Empirique à Partir des Réformes Françaises de 1995 et 2000." *Économie et Statistique* 413 (1): 3–20.
- Castanheira, Micael, Gaëtan Nicodème, and Paola Profeta. 2012. "On the Political Economics of Tax Reforms: Survey and Empirical Assessment." *Internat. Tax and Public Finance* 19 (4): 598–624.

- Chirakijja, Janjala, Thomas Crossley, Melanie Lührmann, and Cormac O’Dea. 2009. “The Stimulus Effect of the 2008 UK Temporary VAT Cut.” Paper presented at the 102nd annual conference of the National Tax Association, Denver, CO, November 12–14.
- Crump, Richard K., V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. 2009. “Dealing with Limited Overlap in Estimation of Average Treatment Effects.” *Biometrika* 96 (1): 187–99.
- Doyle, Joseph J., and Krislert Samphantharak. 2008. “\$2.00 Gas! Studying the Effects of a Gas Tax Moratorium.” *J. Public Econ.* 92 (3): 869–84.
- European Commission. 1999. “Council Directive 1999/85/EC.” *Official J. European Union*, October 22.
- . 2006. “Council Directive 2006/18/EC.” *Official J. European Union*, February 14.
- Evans, William N., Jeanne S. Ringel, and Diana Stech. 1999. “Tobacco Taxes and Public Policy to Discourage Smoking.” *Tax Policy and Econ.* 13:1–55.
- Feldstein, Martin, and Paul Krugman. 1990. “International Trade Effects of Value-Added Taxation.” In *Taxation in the Global Economy*, edited by Assaf Razin and Joel Slemrod, 263–82. Chicago: Univ. Chicago Press.
- Foremny, Dirk, and Nadine Riedel. 2014. “Business Taxes and the Electoral Cycle.” *J. Public Econ.* 115:48–61.
- Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch. 2018. “Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany.” *A.E.R.* 108 (2): 393–418.
- Fullerton, Don, and Gilbert E. Metcalf. 2002. “Tax Incidence.” In *Handbook of Public Economics*, vol. 4, edited by Alan J. Auerback and Martin S. Feldstein, 1787–1872. Amsterdam: Elsevier.
- Gaarder, Ingvil. 2018. “Incidence and Distributional Effects of Value Added Taxes.” *Econ. J.* 129 (618): 853–76.
- Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski. 2011. “Optimal Dynamic Taxes.” Working Paper no. 17642 (December), NBER, Cambridge, MA.
- Hallerberg, Mark, and Carlos Scartascini. 2017. “Explaining Changes in Tax Burdens in Latin America: Do Politics Trump Economics?” *European J. Polit. Econ.* 48:162–79.
- Heckman, James J., Hidehiko Ichimura, and Petra E. Todd. 1997. “Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme.” *Rev. Econ. Studies* 64 (4): 605–54.
- Hines, James, and Mihir Desai. 2005. “Value-Added Taxes and International Trades: The Evidence.” Working paper, Bus. School, Univ. Michigan.
- Imbens, Guido W. 2015. “Matching Methods in Practice: Three Examples.” *J. Human Resources* 50 (2): 373–419.
- Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote. 2001. “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players.” *A.E.R.* 91 (4): 778–94.
- Karadi, Peter, and Adam Reiff. 2019. “Menu Costs, Aggregate Fluctuations, and Large Shocks.” *American Econ. J.: Macroeconomics* 11 (3): 111–46.
- Kopczuk, Wojciech, and Joel Slemrod. 2006. “Putting Firms into Optimal Tax Theory.” *A.E.R.* 96 (2): 130–34.
- Kosonen, Tuomas. 2015. “More and Cheaper Haircuts after VAT Cut? On the Efficiency and Incidence of Service Sector Consumption Taxes.” *J. Public Econ.* 131:87–100.
- Kotlikoff, Laurence, and Lawrence Summers. 1987. “Tax Incidence.” In *Handbook of Public Economics*, vol. 2, edited by Alan J. Auerback and Martin S. Feldstein, 1043–92. Amsterdam: North-Holland.

- Mirrlees, James. 2010. *Dimensions of Tax Design: The Mirrlees Review*. Oxford: Oxford Univ. Press.
- Naritomi, Joana. 2019. "Consumers as Tax Auditors." *A.E.R.* 109 (9): 3031–72.
- Peltzman, Sam. 2000. "Prices Rise Faster Than They Fall." *J.P.E.* 108 (3): 466–502.
- Politi, Ricardo Batista, and Enlison Mattos. 2011. "Ad-Valorem Tax Incidence and After-Tax Price Adjustments: Evidence from Brazilian Basic Basket Food." *Canadian J. Econ.* 44 (4): 1438–70.
- Pomeranz, Dina. 2015. "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax." *A.E.R.* 105 (8): 2539–69.
- Slemrod, Joel, and Christian Gillitzer. 2013. *Tax Systems*. Cambridge, MA: MIT Press.