UCLouvain

Economics School of Louvain

Essays on Commodity Tax Incidence

Valerio Serse

4th September 2020

Members of the Jury

Jean Hindriks (UCLouvain) – Supervisor

Sébastien Van Bellegem (UCLouvain) – President of the Jury

Pierre Dubois (Toulouse School of Economics)

Axel Gautier (ULiège)

William Parienté (UCLouvain)

To my grandparents

Acknowledgements

I want to thank my supervisor, Jean Hindriks, for his constant support throughout this research project. He has encouraged me to follow my research interests and helped me to deal with the many difficulties of a Ph.D. program. I would also like to thank the members of the jury for their careful reading of this work. Their valuable comments and suggestions have substantially improved the quality of this dissertation. I am very grateful to the FNRS for having financed this research project. In this regard, I would like to thank Professors Paul Belleflamme and Frank Verboven, which have sponsored my application to the Ph.D. scholarship. I am also very grateful to the Colruyt Group, the Commission for Electricity and Gas (CREG) in Belgium, and Synergrid. In particular, I would like to thank Dries De Dauw and Bart Van Roost of the Colruyt Group, Laurent Jacquet of the CREG, and Bruno Gouverneur of Synergrid. They gave me the unique opportunity to use their data for my research agenda. Their trust and interest in this research have always pushed me to improve my work. I would like to thank the members of the data analytics team at Colruyt Annelies Bartlema, Lode Nachtergaele, Davy Sannen, Steve Van Straaten, and Wannes Vandesande. I have learned a lot from them, and I am extremely grateful for their patience in helping me to manage and understand raw data. I would also like to thank Natalie Cornelis and Cristine Cobut of the price analytics team at the CREG. Their help in understanding price data and the functioning of the electricity market was crucial to complete the last chapter.

I would like to thank my friends and colleagues at CORE and Economics School of Louvain. They have contributed to this research not only by exchanging remarks and ideas, but also by creating a friendly and stimulating working environment, which I will certainly miss for the years to come. In particular, I would like to thank Fabrizio Ciotti, Andras Gregor, Huan Ha, Leonardo Madio, Hassan Nostratabadi, Erica Pini, and Lorenzo Tondi. I would also like to thank the visitors at the Economics School of Louvain and the numerous participants at conferences and seminars around Europe and the U.S. during these four years of research.

Lastly, I am very grateful to my family for their constant support during these years abroad. They have always pushed me to follow my passions, although our distance was an inevitable price to pay. I feel fortunate to have them on my side.

Table of Contents

Chapter I

Introduction

The economic literature on tax incidence studies how the burden of a tax is distributed across market agents (firms and consumers). The study of tax incidence is notably at the intersection between public economics and industrial organization, and it has interested economists for a very long time. Economic theorists have provided flexible frameworks to predict the impact of commodity taxes on the market outcome. In particular, they have shown that tax incidence depends on three fundamental characteristics of the market, which are the elasticity of demand and supply and the type of competition. These features can vary substantially across markets, geographical areas, and time, but they can nonetheless be measured empirically. Thanks to the increasing availability of market data, many economists have recently focused on the empirical investigation of tax incidence in a wide variety of markets. This dissertation contributes to this literature by collecting three empirical studies of commodity tax incidence in the markets for sodas, spirits, and residential electricity.

The starting point for the study of tax incidence is that the agent who actually pays the tax may not be the one that is formally liable for the tax payment. The *economic incidence* of a tax identifies the agents that bear the burden of taxation. In contrast, the *legal incidence* of the tax refers to the agents having the legal obligation to remit the tax to the fiscal authority. Interestingly, economic and legal incidence rarely coincide. Let us take, for instance, the case of excise duties on alcohol. These taxes are typically levied on producers and are based on the amount of alcohol contained in a product. However, once these taxes are implemented, the retail price of alcoholic beverages is likely to increase, thus sharing the burden of the tax between consumers and producers. An additional element of interest is how the burden of the tax is distributed across different types of agents. While commodity taxes are usually uniform across agents, their behavioral response to these taxes can be highly heterogeneous. As a result, some consumers (or producers) can bear a higher tax burden as compared to others.

To understand the importance of studying tax incidence, it is fundamental to remind why governments tax commodities in the first place. One of the primary rationales for commodity taxation is the necessity of governments to raise fiscal revenues. However, taxing commodities imposes a cost on economic agents. This leads to the question of how to set the optimal commodity tax rates that allows raising revenues without affecting too much market efficiency.

An initial solution to this problem was given by the seminal work of Ramsey (1927), which proposes an optimal tax rule (i.e., the *Ramsey rule*) to maximize welfare subject to a revenue constraint. The rule relates directly optimal tax rates to the elasticity of demand and supply. The lower the sum of demand and supply elasticity, the higher the tax rate.

Another rationale for commodity taxation is to correct inefficiencies due to *externalities* in the production or consumption of a product. In this case, taxes are not set to minimize the welfare loss but to increase economic efficiency. Taxes having such a goal are also known as *Pigouvian taxes,* named after the work of Pigou (1920). The optimal tax to address an externality must be equal to the marginal external damage from the production or consumption of the good. This optimality, however, it is difficult to achieve in practice for at least two reasons. First, to achieve the first-best outcome, these taxes must usually vary across both commodities and market agents. Second, the exact identification of the marginal damage is often a very cumbersome empirical exercise. Importantly, Pigouvian taxes are effective in internalizing an externality as long as agents respond to them. If demand for the product is inelastic, then such a tax will raise fiscal revenues, but it will not be effective in reducing demand and quantity at equilibrium.

The advent of behavioral economics has also opened up to the possibility of using commodity taxation as a tool to address *limited rationality* in consumer choice (Thaler, 1980). Alongside the concept of externalities, it is now that of *internalities*. Internalities are costs that individuals impose on themselves but are not taken into account when making choices. For instance, recent works focused on the optimal design of *sin taxes* (i.e., taxes on harmful goods), considering both internalities and redistributive concerns (see Gruber & Kőszegi, 2004; O'Donoghue & Rabin, 2006; Griffith et al., 2018; Allcott et al., 2019). In general, this literature highlights that optimal corrective taxes depend on the empirical distribution of internalities and how they correlate with the shape of demand and income distribution. Similar to Pigouvian taxes, corrective taxes are effective in addressing internalities as long as biased consumers are sensitive to price changes.

Commodity taxation does not only affect consumers. Taxes do also represent a shock for firms. A tax on a product drives a wedge between the consumer and producer price, which are normally equal in the absence of taxation. In perfect competition, with constant marginal cost, any tax is entirely shifted to consumer the price, leaving the producer price unchanged. With upward-sloping marginal cost, however, the tax burden is shared between producers and consumers and depends on the relative elasticity of demand and supply (Fullerton & Metcalf, 2002). It is the most inelastic side of the market that bears the tax the most. Tax incidence under perfect competition is quite straightforward and facilitates the analysis of optimal taxes. Nevertheless, in real markets, perfect competition may be more an exception rather than the norm.

Tax incidence in imperfectly competitive markets is more complex and makes optimal tax design also more challenging. Under imperfect competition, firms enjoy some market power and are able to price above the marginal cost. The change in consumer price after the tax may not be equivalent to that one under perfect competition. In some cases, the increase in consumer price can also exceed the amount of the tax (i.e., *over-shifting*), a result that is not possible in the competitive case. A new fundamental variable to understand tax incidence in imperfectly competitive markets is the curvature of demand. For firms facing a convex demand curve (demand becomes increasingly steep as price increases), it is easier to pass the tax more on consumers (Weyl & Fabinger, 2013). Furthermore, imperfect competition complicates the optimal tax design, as the welfare loss due to market power must also be taken into account.

A sound understanding of tax incidence is key in designing and implementing optimal commodity taxes. Economic theory provides a rigorous framework to analyze the impact of a tax on the market outcome. The elasticity of demand and supply and the type of market structure are the three main determinants of tax incidence. These characteristics, however, must be measured empirically and can vary substantially across markets and geographical areas. Over the last two decades, there has been a large development in the empirical literature on tax incidence, which was facilitated by the increasing availability of detailed market data. This literature can usually follow two empirical approaches: a *reduced-form* approach, which focuses on identifying the statistical relationship between the price and the tax by exploiting the implementation of exogenous tax reforms (e.g., Kenkel, 2005; Doyle & Samphantharak, 2008; Harding et al., 2012); or a *structural* approach, which specifies a model of economic behavior that is then estimated on market data and allows to simulate the impact of a tax on the market outcome (e.g., Bonnet & Réquillart, 2013; Dubois et al., 2020; Miravete et al., 2018).[1]

This dissertation contributes to the empirical literature on tax incidence by collecting three articles focusing on the impact of commodity taxation using different types of data and statistical methods. Chapter II studies the incidence of sugar taxes on sodas using a structural approach that allows estimating demand and simulating the impact of taxation on different consumer types. Chapter III analyzes the incidence of an alcohol tax reform implemented in

[1] See Chetty (2009) for an interesting discussion about the use of these two approaches in public economics.

Belgium in 2015, focusing on the spatial variation of tax incidence across stores and geographical areas. Chapter IV studies the impact of a temporary VAT reform on residential electricity price and demand. The results of these studies can provide insights into the behavior of market agents and support the design of optimal tax policies. Throughout this dissertation, a particular focus is given to the heterogeneity of consumers and firms. The main finding is that accounting for heterogeneity in consumer preferences and firms' behavior is crucial to correctly identify and measure tax incidence.

A novelty of this dissertation is the use of unique datasets that allows undertaking a robust analysis of tax incidence. Chapter II exploits a retail chain *scanner data* that tracks household shopping decisions by using their loyalty cards. The dataset includes information about almost half a million Belgian households (out of a total of 4.5 million households) doing shopping over a period of two and a half years, for a total of 9.7 million shopping trips. All the next three chapters make use of *posted price* data instead of the conventional data price typically used in the literature (e.g., Nielsen measured prices). Posted prices tend to be less sensitive to local and cyclical shocks (Coibion et al., 2015) and are not dependent on measurement errors due to loyalty cards (Einav et al., 2010). The next two chapters use posted price data from a major Belgian supermarket chain, which varies at the product-day-store level. Chapter IV uses the residential (and business) electricity price that is posted in the monthly electricity bill of each service provider.

Each article in this dissertation has a specific contribution to the existing literature on tax incidence. Chapter II investigates whether sugar taxes are effective in targeting heavy sugar consumers. This research question has been recently addressed by various economists, but results are still mixed. Bonnet & Requillart (2018) show that soda consumption increases with the Body Mass Index (BMI) in both adults and children. Their findings can suggest that taxes on sodas can be effective in targeting unhealthy individuals, but they do not document any heterogeneity in price sensitivity across consumers. Dubois et al. (2020) study the individual demand for drinks on-the-go in the UK. They find that heavy sugar consumers prefer sugary drinks and are relatively insensitive to price, thus suggesting that a sugar tax is not well targeted at individuals over-consuming sugary drinks. The article in chapter II extends this literature by studying household demand for bottled colas using a large supermarket dataset from a Belgian retail chain. Like Bonnet & Requillart (2018) and Dubois et al. (2020), it is estimated a *discrete choice model* of demand, which accounts for both unobserved and observed heterogeneity in consumer preferences (McFadden, 1974; Train, 2009), and that allows controlling for dynamics in cola demand. The model estimates then allow performing tax policy simulations to study tax incidence across different household types.

The results show that households consuming more sugary products tend to prefer sugary to diet colas while they are less sensitive to cola prices. However, low-income households have a greater preference for sugary colas and are relatively more price-sensitive than high-income households. These findings suggest that sugar taxes would be less effective among heavy sugar consumers, but they would be quite effective among the poor. This result is in line with Dubois et al. (2020), which focused on the market for soda on-the-go and individual level demand in the UK. Therefore, the results of chapter II extend their findings to the market for larger bottle size and household level demand in another geographical context. This chapter also provides novel evidence of demand dynamics in the context of sin good markets. Indeed, because of state dependence in cola choice, the impact of sugar taxes would be larger in the long-run. Tax policy simulations also show that a *specific* tax on sugar may be preferred to an *ad-valorem* tax on both corrective and equity grounds. This is because an ad-valorem tax would imply a substitution toward cheap sugary brands, especially among low-income households and heavy sugar consumers.

Chapter III studies the incidence of an alcohol tax hike that occurred in Belgium in 2015. Many works investigated tax incidence for alcoholic beverages (e.g., Kenkel, 2005; Carbonnier, 2013; Ally et al., 2014). However, little attention has been given to the spatial heterogeneity in tax incidence for homogeneous products. This issue was addressed in the context of cigarettes (Harding et al., 2012) and fuel taxes (Doyle & Samphantharak, 2008), but not for alcohol taxation. Understanding the spatial heterogeneity in alcohol tax incidence can have important implications for alcohol control policies. In fact, if tax incidence varies over space, these policies risk affecting households differently depending on where they live. The article in chapter III focusses on estimating and explaining spatial heterogeneity in alcohol tax incidence. This is done by analyzing the (posted) retail price of six major brands of spirits, using a *difference-in-differences* method. The estimation uses the retail price of these brands sold in a major supermarket chain and uses the retail prices of the same brands sold in France by the same supermarket chain as a control group. This chain has over 400 stores in France and Belgium. Since there is information on each store location, it is then possible to test for and explain spatial heterogeneity in tax incidence.

The results show a substantial variation in the tax incidence across stores for homogeneous products. Interestingly, these variations are strongly related to the intensity of local competition and, to a lesser extent, to the proximity to the borders (mainly with Luxembourg, which is the low-price country). The tax hike was quickly passed through during the first month of implementation, and it was mostly over-shifted. Furthermore, we also find that both the border and the competition effects are not instantaneous but arise several months after the tax reform. From a public health perspective, these findings suggest that the health benefits associated with the tax reform could have a differential impact on Belgian households according to where they live. This hypothesis is supported further by analyzing the evolution of spirit sales in the stores considered before and after the reform, which shows a very heterogeneous variation in demand over Belgian provinces. The analysis of spirit sales also provides evidence of stockpiling before the tax reform and a substantial rise in spirit sales in Luxembourg, indicating effective cross-border shopping by Belgian consumers.

Chapter IV analyzes the impact of a temporary VAT reform in the Belgian electricity market, which occurred in the period between April 2014 and September 2015. The Belgian government temporarily reduced the VAT rate on electricity prices from 21% to 6%. This article studies how such a VAT cut affected both the price and demand for electricity. The incidence of VAT in Europe has been investigated in many markets (e.g., Carbonier, 2007; Kosonen, 2015; Benedek et al., 2019), but not in the one of residential electricity. This literature typically finds that VAT reductions are not entirely passed to consumer prices. Like Chapter III, the VAT incidence to electricity prices is estimated by means of a difference-in-differences analysis, where it is used the electricity price paid by firms (not subject to VAT) as a control group. In contrast with the existing literature, the results show that the VAT cut was fully shifted to electricity prices so that households could entirely benefit from the tax reduction.

The VAT reform also provides a useful exogenous shock to estimate the price elasticity of residential demand for electricity. Many works studied the demand elasticity for electricity (see, for instance, Labandeira et al., 2017, for a review), although very few papers are based on *quasi-experimental* (exogenous) variation in prices. Ito (2014) estimates the demand response to price variations induced by a discontinuity in electricity service areas in California. While Deryugina et al., (2020) study the elasticity of residential electricity demand following an exogenous drop in electricity prices due to a municipal aggregation policy in Illinois. The impact of the VAT cut on demand and the resulting demand elasticity are estimated by analyzing the monthly variation in consumption of electricity at the network distributor level, controlling for changes in other determinants of energy use during the period. The results show that the VAT cut generated roughly a 2% increase in demand for electricity. By exploiting different sources of price variation, the analysis reveals a price elasticity of residential demand for electricity between -0.09 and -0.17, which is in the range of other elasticities based on quasi-experimental variation in price (Ito, 2014; Deryugina et al., 2020). Interestingly, results also show that households reacted quickly and symmetrically (albeit in a different direction) to the VAT cut and the successive VAT hike.

The next three chapters are structured as follows. They start with a general introduction of the topic to be analyzed. It is then presented a review of the relevant economic literature, together with a discussion of the contribution of the article. Data and methods are presented in two separate sections. Each article then shows the results of the empirical analysis and some robustness checks to validate them. The articles finish with a brief conclusion, which summarizes the analysis and highlights the possible implications for tax policy.

References

Allcott, H., Lockwood, B. B., & Taubinsky, D. (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *The Quarterly Journal of Economics*, *134*(3), 1557–1626.

Ally, A. K., Meng, Y., Chakraborty, R., Dobson, P. W., Seaton, J. S., Holmes, J., ... & Meier, P. S. (2014). Alcohol tax pass-through across the product and price range: do retailers treat cheap alcohol differently?. *Addiction*, *109*(12), 1994-2002.

Benedek, D., De Mooij, R. A., Keen, M., & Wingender, P. (2019). Varieties of VAT pass through. *International Tax and Public Finance*, 1-41.

Bonnet, C., & Réquillart, V. (2013). Tax incidence with strategic firms in the soft drink market. *Journal of Public Economics*, *106*, 77-88.

Bonnet, C., & Réquillart, V. (2018). *The effects of taxation on the individual consumption of sugar-sweetened beverages* (No. 16-638). Toulouse School of Economics (TSE).

Carbonnier, C. (2007). Who pays sales taxes? Evidence from French VAT reforms, 1987–1999. *Journal of Public Economics*, *91*(5-6), 1219-1229.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, *1*(1), 451-488.

Coibion, O., Gorodnichenko, Y., & Hong, G. H. (2015). The cyclicality of sales, regular and effective prices: Business cycle and policy implications. *American Economic Review*, *105*(3), 993-1029.

Deryugina, T., MacKay, A., & Reif, J. (2020). The long-run dynamics of electricity demand: Evidence from municipal aggregation. *American Economic Journal: Applied Economics*, *12*(1), 86-114.

Diamond, P. A., & Mirrlees, J. A. (1971). Optimal taxation and public production I: Production efficiency. *The American Economic Review*, *61*(1), 8-27.

Doyle Jr, J. J., & Samphantharak, K. (2008). $2.00 Gas! Studying the effects of a gas tax moratorium. *Journal of public economics*, *92*(3-4), 869-884.

Dubois, P., Griffith, R., & O'Connell, M. (2020). *How well targeted are soda taxes?* (No. 17-868). Toulouse School of Economics (TSE).

Einav, L., Leibtag, E., & Nevo, A. (2010). Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?. *QME*, *8*(2), 207-239.

Fullerton, D., & Metcalf, G. E. (2002). Tax incidence. *Handbook of public economics*, *4*, 1787-1872.

Griffith, R., O'Connell, M. & Smith, K. (2018). Corrective taxation and internalities of food consumption. *CESifo Economic Studies*, *64,* 1–14.

Gruber, J., & Kőszegi, B. (2004). Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes. *Journal of Public Economics*, *88*(9-10), 1959-1987.

Harding, M., Leibtag, E., & Lovenheim, M. F. (2012). The heterogeneous geographic and socioeconomic incidence of cigarette taxes: Evidence from Nielsen Homescan Data. *American Economic Journal: Economic Policy*, *4*(4), 169-198.

Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review*, *104*(2), 537-63.

Kenkel, D. S. (2005). Are alcohol tax hikes fully passed through to prices? Evidence from Alaska. *American Economic Review*, *95*(2), 273-277

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

Labandeira, X., Labeaga, J. M., & López-Otero, X. (2017). A meta-analysis on the price elasticity of energy demand. *Energy Policy*, *102*, 549-568.

McFadden, D. (1974), Conditional logit analysis of qualitative choice behavior. P. Zarembka (Ed.), *Frontiers in Econometrics*, Academic Press, New York, 105–142.

Miravete, E. J., Moral, M. J., & Thurk, J. (2018). Fuel taxation, emissions policy, and competitive advantage in the diffusion of European diesel automobiles. *The RAND Journal of Economics*, *49*(3), 504-540.

O'Donoghue, T., & Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, *90*(10-11), 1825-1849.

Pigou, A. C. (1920). *The Economics of Welfare*. London: Macmillan.

Ramsey, F. P. (1927). A Contribution to the Theory of Taxation. *The Economic Journal*, *37*(145), 47-61.

Thaler, R. (1980). Toward a positive theory of consumer choice. *Journal of economic behavior & organization*, *1*(1), 39-60.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Weyl, E. G., & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528-583.

Chapter II

Taste Heterogeneity and State Dependence in Cola Choice: How do Sugar Taxes affect Households' Demand?

Author: Valerio Serse[†]

Abstract: A correct understanding of *taste heterogeneity* across households is crucial to assess the impact of sugar taxes in the population. Taste heterogeneity, however, can be easily confounded with other sources of choice inertia, such as *state dependence* in product choice. Distinguishing between these two behaviors empirically can have important policy implications. Taste heterogeneity implies that a sugar tax has different impact across consumers, while state dependence implies that demand is dynamic and so the impact of a tax increases with time. This paper estimates a dynamic multinomial logit model of cola demand on a novel supermarket scanner dataset to study both preference heterogeneity across households and state dependence in cola choice. The model estimates allow exploring how sugar taxes can affect the shopping decision of households with different levels of income and sugar consumption, while controlling for possible dynamics in cola demand. The results show that households have very heterogenous taste for colas and exhibit state dependence in product choice. Tax policy simulations show that a sugar tax would be less effective in reducing sugary cola demand among the targeted population of heavy sugar consumers. This policy, however, would be more effective among low-income households. Furthermore, because of state dependence in cola choice, this policy would be more effective in the long-run, as consumers can take time to fully react to a price increase.

JEL No: D12, H31, I18, Q18.

Keywords: preference heterogeneity, state dependence, sugar taxes, discrete choice models.

[†] CORE, Université catholique de Louvain, FNRS-FRESH grant holder, e-mail: valerio.serse@uclouvain.be.
Acknowledgements: I am very grateful to my advisor, Jean Hindriks, for helpful comments and discussions. I also would like to thank Pierre Dubois, Liran Einav, Johannes Johnen, Leonardo Madio, Hassan Nosratabadi, Martin O'Connell, Sébastien Van Bellegem, Frank Verboven and participants at the DSSG conference in Chicago, the ECORES summer school in LLN, the EC^2 conference in Rome, the RGS doctoral conference in Bochum, the WIPE workshop in Reus, the PET conference in Strasbourg, the PhD seminar at AMSE in Marseille and various seminars at the UCLouvain. The financial support from the FNRS is gratefully acknowledged.

2.1 Introduction

The obesity rate has substantially increased during the last decades (WHO, 2017), generating high costs in terms of healthcare and foregone productivity (Dobbs et al., 2014). Excessive sugar consumption is often blamed as one of the main drivers of this negative trend (Te Morenga et al., 2013). In this regard, some countries have already taken actions by implementing taxes aimed to discourage the excessive consumption of sugar-sweetened beverages (SSB). Assessing the effectiveness of these policies, however, requires a careful evaluation of consumers' reaction to price changes. While taxing tobacco affects mostly smokers, taxing SSB can affect various types of households, even those with low sugar consumption and a low health risk. It is therefore fundamental to understand whether sugar taxes will affect the targeted population of heavy sugar consumers and how the tax burden will fall on different population groups.

This paper studies the impact of sugar taxes on cola demand across different types of households. The aim is to understand how sugar taxes affect the probability of switching from a sugary to a diet cola (i.e., sugar-free) across households with different levels of income and sugar consumption. The analysis consists of estimating a discrete choice model of cola demand on a novel supermarket scanner dataset from a major retail chain in Belgium. The dataset includes information about the socioeconomic status and dietary habits of each household unit, which is tracked over time by the use of loyalty cards. The model estimates are then used to perform tax policy simulations, which allows studying the impact of different levels of taxation on sugary cola demand. These simulations are also used to compare the effectiveness of ad-valorem sugar taxes (i.e., taxes based on the product price) to *specific* sugar taxes[2] (i.e., taxes based on sugar content). The findings of this paper provide useful empirical insights for the optimal design of sugar taxes by focusing on their effectiveness in reducing demand among the targeted population (i.e., heavy sugar consumers) and their redistributive concerns.

The demand model accounts for both *preference heterogeneity* and *state dependence* in product choice. Preference heterogeneity means that households have some persistent differences in preferences for colas. State dependence instead means that households have a higher probability of choosing a product they have purchased in the recent past. These two behaviors have important policy implications. Accounting for preference heterogeneity in the demand model is crucial as it gives a measure of how different types of households, such as those with a high-

[2] Sometimes referred as excise sugar taxes.

sugar diet, would react to sugar taxes. Accounting for state dependence is also important as it implies that household demand is dynamic, and so the impact of sugar taxes on demand will tend to increase over time. State dependence and taste heterogeneity are not mutually exclusive, and there is a risk of empirically confound one with the other. This is because they can generate a similar sequence of choices with the consumer that keeps on choosing the same product over time. Hence, allowing simultaneously for both behaviors in the demand model is necessary for their correct identification (Heckman, 1978).

The model is estimated on a novel set of supermarket scanner data provided by a major group of Belgian retailers. The dataset is composed of 434,475 households that purchased at least five times one or more bottles of cola over two and a half years, for a total of 9.7 million shopping trips. Colas are an interesting market to analyze for at least two reasons. First, they make up the largest share of SSB, which are often considered as a possible target for public health taxes. Second, each of these products has a perfect healthier substitute, as most brands also have a diet version. Since this retail chain is a local price follower, there is substantial variation in cola prices both across stores and over time. Price data consists of the daily posted price specific to each cola in a given store. The dataset also includes information about the presence of quantity rebates and whether the product was advertised in the supermarket's weekly folder. Furthermore, there is data about which product was available in the stores on a given day. This provides complete information on the choice set faced by each household during any shopping trip.

The cola choice model accommodates random taste variation in product attributes and hence allows cross-price elasticities to depend on the similarity between products (Train, 2009). Households are allowed to have a different price sensitivity and heterogeneous preferences for sugary colas and bottle size. As the model can capture heterogeneity in price sensitivity across households, it also allows approximating the curvature of product demand. Which is an essential variable in predicting tax pass-through in imperfectly competitive markets (Weyl & Fabinger, 2013). The unobserved heterogeneity in taste across households is specified as a parametric distribution of the random taste coefficients. In order to understand how tastes change with observed household characteristics, the means of these distributions are also allowed to vary with dietary habits and household demographics. State dependence is tested with the inclusion of the lagged choice variable in the utility specification while controlling for persistent heterogeneity in household preferences.

By exploiting the time dimension of the data, this paper studies preference heterogeneity across households and disentangle it from state dependence in cola choice. The identification of taste heterogeneity against state dependence is possible thanks to the frequent temporary price discounts and variations in the consumer choice set due to changes in product availability over time. The following example may give an idea of how these temporary market shocks are exploited for identification purposes. Imagine that there are only two brands, A and B, and that a household prefers brand A whenever it is available. One day brand A is not available, so the household chooses brand B. In the next shopping trip, brand A is back on the market. In the absence of state dependence, the household should revert to brand A as it is again available. If it keeps choosing brand B, however, the choice will exhibit state dependence. This is because choosing B in the previous shopping trip (when A was not available) leads the household to choose B to A in the next shopping occasion (when A is again available).[3]

The results of this paper show that households have very heterogeneous preferences for colas and are state-dependent in cola choice. Accounting for state dependence reduces the degree of unobserved taste heterogeneity by one third. This suggests that ignoring state dependence in demand estimation could overstates the extent of taste heterogeneity across households. This finding is robust to possible misspecifications of taste heterogeneity and autocorrelation in the error terms over time. The model estimates show that heavy sugar consumers tend to prefer sugary to diet colas, while they are less sensitive to cola prices. This finding suggests that taxing sugary colas would be less effective in reducing demand among the targeted population and that the tax burden would mostly fall on these consumers. However, the results also show that low-income households are likely to experience larger health gains compared to richer households. This is because they have a greater preference for sugary colas and are relatively more price-sensitive than high-income households. Their higher responsiveness to sugar taxes also helps them reduce their tax burden in monetary terms.

The model estimates allow performing tax policy simulations, which are then used to compare the impact of different tax schemes on sugary cola demand. This analysis shows that specific taxes targeting sugar content should be preferred to ad-valorem taxes on sugary colas on both corrective and equity grounds. The main reason for this finding is that ad-valorem taxes would entail a larger substitution towards cheaper sugary brands. This effect is mostly concentrated

[3] A similar reasoning can be done for temporary price discounts. Imagine a temporary price discount on brand B leading the household to switch from A to B. Once the price discount is over, the household should revert to its favorite brand A. Yet, if the household is state dependent, brand B may still be chosen against A, even though the discount is terminated.

among low-income households and heavy sugar consumers. Finally, because of state dependence in cola choice, households need some time to fully adjust their demand to a sugar tax. Hence, these policies would tend to be more effective in the long-run. This effect has also the consequence of reducing the government revenues from sugar taxes over time, as less consumers would purchase a sugary cola.

The paper is structured as follows. Section 2.2 reviews some of the most recent findings in the economic literature focusing on consumer demand estimation and junk food taxation. Section 2.3 describes the scanner data used for the estimation. Section 2.4 presents the cola demand model and its empirical specifications. Section 2.5 discusses the parameter estimates and shows the results of various tax policy simulations. In section 2.6, it is presented an extended version of the model, which accounts for both observed and unobserved heterogeneity in household preferences. Section 2.7 shows the results of various robustness checks. Section 2.8 concludes the paper.

2.2 Background Literature

A substantial body of economic literature evaluates the impact of commodity taxes by means of consumer demand estimation. This approach allows estimating ex-ante the consumer response to a given change in tax policy. This exercise can shed a light on various aspects related to this type of policy. Such as, for instance, their distributional effect, their capacity to generate fiscal revenues or their effectiveness in reducing the consumption of unhealthy products.

It is a common practice to adopt models of consumer demand that are both consistent with economic theory and flexible enough to be estimated on commonly available purchase data. One of the most successful models in this fashion is the Almost Linear Demand System (AIDS) proposed by Deaton & Muellbauer (1980), which is suitable for the most common forms of consumer survey data. Various studies use this approach to estimate the impact of taxes on unhealthy food demand (see, for instance, Jensen & Smed, 2007; Chouinard et al., 2007; Nordström & Thunström, 2009; Allais et al., 2010; Sharma et al., 2014; Harding & Lovenheim, 2017). The main limitation of this class of models, however, is that it is based on a representative household, and hence, it does not focus on preference heterogeneity across consumers.

Understanding consumer heterogeneity plays a crucial role in the optimal design of sin good taxes. This is because policymakers may be primarily interested in the behavioral response of the targeted population. From a public health perspective, the targeted population should be those individuals overconsuming junk food products. Hence, it is important to understand how this particular group of individuals responds to price changes. Understanding preference heterogeneity is also crucial for the optimal tax design in an economic welfare framework. Some theoretical studies focused on the optimal design of sin good taxes, considering both market efficiency and redistributive concerns (see, for instance, Gruber & Kőszegi, 2004; O'Donoghue & Rabin, 2006). More recently, Griffith et al. (2018) study optimal corrective taxes with internalities in food consumption. They show that optimal corrective taxes depend on the empirical distribution of these internalities and how they correlate with the shape of demand and income distribution. In particular, they show that junk food taxes could be both efficient and progressive if low-income consumers are very sensitive to prices, have a higher marginal utility of income, consume more unhealthy food, and suffer more from internalities.

Similar results are obtained by Allcott et al. (2019), who study the optimal sin good tax with both internalities and externalities in consumption accounting for redistributive concerns and tax revenue redistribution. They provide an optimal sin tax formula that can be directly computed using three sufficient statistics. Such as the consumer elasticities related to prices and income, a money-metric measure of consumer bias, and the progressivity of bias correction through taxation. By calibrating their model on Nielsen Homescan data, they find that the optimal soda tax in the US should be between 1 and 2.1 cents per ounce. Their estimate of price and income elasticities is computed by regressing the natural log of SSB purchase on price and income. They address the possible simultaneity bias in demand estimation through an instrumental variable approach. However, they restrict heterogeneity in soda demand by allowing price elasticities to vary only with the income level.

A common way to study preference heterogeneity in consumer demand estimation is to use Multinomial Logit (MNL) models of demand, such as those in Berry (1994), Berry et al. (1995), and Nevo (2000). These models are based on random utility theory and can be estimated on either aggregate or consumer-level data. Contrary to the AIDS, which can describe the allocation of expenditure across different product categories, the MNL is often adopted to study the consumer choice within a single category of products. This class of models sees the consumer as choosing across different goods the one that delivers the highest utility of consumption, which depends on prices and the consumer preference for product attributes. These models can deliver very realistic substitution patterns by allowing the taste for different product attributes to vary randomly across consumers. The resulting price elasticities of demand will depend on the similarity between products in terms of their characteristics. However, in contrast to the AIDS, the MNL does not allow any product in the demand system to be a complement to any other. Thus, the choice set should be constructed carefully by including only the products that are closely substitutable with each other.

SSB are often considered as a possible target for sugar taxes given the established link between their consumption and body weight gain (Malik et al., 2013). Lately, various authors used MNL models of demand to study the potential impact of sugar taxes on the demand for these products. Bonnet & Requillart (2018) estimate a random coefficient MNL on a panel data of French households, accounting for heterogeneity in beverage consumption both within and across households. They show significant differences in the consumption of SSB across individuals. In particular, they find that SSB consumption increases with the Body Mass Index (BMI) in both adults and children. This finding suggests that taxes on SSB can be justified on public health grounds and internality motives. However, they do not document any heterogeneity in price sensitivity across consumers. Wang (2015) proposes a MNL model of soda demand that accounts for both preference heterogeneity and product stockpiling. She shows that ignoring this behavior leads to a substantial overestimation of the price elasticity of demand for SSB. According to her estimates, this policy is unlikely to reduce soda consumption substantially, but it could be quite effective in raising fiscal revenues.

Dubois et al. (2020) circumvent the issue of soda stockpiling by studying the impact of sugar taxes on the demand for on-the-go drinks (which are less likely to be stored). They develop a nonparametric MNL model and estimate it on a panel of UK individuals that includes information about their purchases at the product level. Their model allows each individual to have a specific taste coefficient for the price, soda, and sugary drinks, which are then correlated with their demographics and dietary habits. They find that heavy sugar consumers prefer sugary drinks and are relatively insensitive to price. While young and poor individuals, who also have a strong preference for sugary drinks, tend to be more price-sensitive. They also show that a tax on sugary sodas will achieve a larger reduction in sugar consumption than a more general soda tax since it leads to a greater substitution from sugary to diet sodas. However, heavy sugar consumers would be less likely to switch to a diet soda compared to other consumers. Their findings thus imply that a sugar tax is not well targeted at those individuals over-consuming sugary drinks.

Although these studies account for preference heterogeneity across consumers, they do not allow for other possible sources of choice inertia in soda demand. Heckman (1978) distinguished two types of brand inertia. One is the so-called *structural state dependence*, which means that the history of past purchases directly affects the current choice behavior. While the other is the *spurious state dependence*, which means instead that consumers have some persistent propensity in making certain decisions (e.g., taste heterogeneity). Importantly, these two sources of choice inertia are not mutually exclusive, and risk being empirically confounded with each other. Researchers have extensively documented persistence in product choice due to state dependence using consumer panel data (see Allenby & Lenk, 1994; Keane, 1997; Seetharaman et al., 1999; Erdem et al., 2008; Dubé et al., 2010). Some studies have also found that ignoring state dependence can overestimate the extent of preference heterogeneity across households (Allenby & Lenk, 1994; Keane, 1997). However, none of these works focused on testing state dependence in the context of SSB demand. This literature usually tests for state dependence by including the lagged choice variable in the consumer utility specification, while controlling for persistent heterogeneity in taste across consumers. Such behavior has important implications for corrective tax policies, as the reduction in demand due to a price hike tends to increase over time. As a result, static demand models can underestimate the long-run impact of sugar taxes on SSB demand.

This paper contributes to the economic literature on sugar taxes by estimating a MNL model of cola demand that accounts for both preference heterogeneity and state dependence in product choice. The dataset contains information about household demographics, such as income, educational level, age, and family composition. The 434,475 households included in the dataset represent 10% of all Belgian households. Compared to previous studies, the size of this sample permits exploring in detail how household tastes are distributed across the population. In particular, this allows understanding both observed (through household demographics) and unobserved heterogeneity in price sensitivity, taste for sugar, and package size. The information about households' dietary habits permits studying how price sensitivity and taste for sugar vary across households with different levels of sugar consumption. This analysis allows assessing how a sugar tax can affect different types of households, thus providing useful empirical insights for an optimal sugar tax design.

In contrast with the existing literature, this paper uses posted price data that varies at the day-brand-store level. Posted prices are different from the conventional scanner price data used in the literature (e.g., Nielsen measured prices) as they are not conditional on purchase and thus less sensitive to local and cyclical shocks (Coibion et al., 2015). They are not dependent neither on measurement errors due to loyalty cards (Einav et al., 2010). The high frequency of daily price data at the store level provides a reliable measure of the actual price faced by the household in every choice occasion. This feature overcomes the possible bias arising with higher levels of temporal and spatial price aggregation (Levin et al., 2017). By having information about the daily product availability in each store, the model also allows controlling for choice set heterogeneity across different shopping trips.[4] An issue that is typically ignored in the estimation of SSB demand.[5] The model does not consider any outside option. This approach allows focusing on the purchase probabilities of each brand conditional on buying a cola, thus circumventing the potential issues arising from censored demand, such as stockpiling of storable goods (Hendel & Nevo, 2006; Wang, 2015).[6] To validate the robustness of this approach, it is also shown that including the purchase of other sodas as outside option does not change much the demand parameters.

The results show that sugar taxes would be less effective among households overconsuming sugary products, but they would be quite effective among the poor. This finding is in line with Dubois et al. (2020), which focused on the market for soda on-the-go and individual level demand in the UK. This paper, therefore, extends their findings to the market for larger bottle size and household level demand in Belgium. This work also contributes to the empirical literature on state dependence by providing novel evidence about the existence of such behavior in the context of sin good markets. State dependence implies that the impact of sugar taxes on cola demand is not only heterogenous across households, but it also tends to increase with time. Hence, these policies would be more effective in reducing demand for sugary beverages in the long-run, although they would generate less tax revenues over time.

Lastly, tax policy simulations show that a specific tax on sugar may be preferred to an ad-valorem tax on both corrective and equity grounds. This is because an ad-valorem tax would imply a substitution toward cheap sugary brands, especially among low-income households. This result is consistent with Alvarado et al. (2019), who study the impact of the 10% ad-valorem tax on SSB implemented in Barbados in 2015. They report that while sales of the most expensive brands have declined by 7.2%, those of mid-range SSB have instead increased by 6.5%, suggesting brand down-switching due to the implementation of the ad-valorem tax.

[4] See Goeree (2008) on how ignoring choice set heterogeneity can lead to biased estimates of price elasticities in the context of PC demand.

[5] Dubois et al. (2020) allow the consumer choice set to vary across retailer type per year. Whereas this study allows for choice set heterogeneity at the day-store level.

[6] Contrary to Dubois et al. (2020), who study the market of soda on-the-go, this paper focuses on the demand for colas in bottle (above 1 liter), which are more likely to be stored.

2.3 The Data

The data used in this work are provided by a major group of food retailers in Belgium. This group controls more than 300 local stores and has a market share of 33% (in 2017) in the country. The dataset contains information on all purchases made by households in the supermarkets of the group. Thanks to the use of loyalty cards, it is possible to follow the purchase history of each household unit over time. This data allows estimating both persistent heterogeneity in taste across consumers and state dependence in product choice. Having information about the full shopping basket for each household, it is also possible to compute the share of sugary-dense products on the overall food purchases. This variable is used as a proxy for the household's habits in sugar consumption.

The dataset is constructed as follows. There are three different data sources: (i) the purchase data, (ii) household demographics, and (iii) the store data. Purchase data comes from the scanning of each item at the pay desk. The dataset includes information about every product purchased during any shopping trip at this retail chain. Contrary to most scanner data, this information is at the single shopping trip level. Thus, it is not a weekly or monthly aggregation of households' purchases. The purchase history of each household is collected through the use of loyalty card in the payment transaction. Household demographics are collected at the residential level by an external company through the household's address. Store data are automatically collected daily by the retail chain. This data includes information about the market environment of the store, such as prices, the set of available products, and the presence of quantity rebates. These three sources of data are then merged to recreate the original choice set faced by the household during the day of purchase, together with the information about the products and the household characteristics.

This paper focusses on the demand for cola products, which represents two-thirds of their total SSB purchases. In the period going from August 2015 to January 2018, 434,475 households purchased at least five times one or more bottles of cola in these stores. For each of these households, there is information about their demographics either at the district of residence or at the residential address level.[7] Sugar consumption is measured as the share of sugary-dense products on the overall food and beverage expenditures for each year.[8] However, inferring eating habits for households making sporadic purchases in these stores could be misleading.

[7] A district in Belgium is a statistical subdivision of the commune or municipality. The total number of districts in Belgium is 19,782. This level of aggregation allows having a good approximation of the actual households' demographics.

[8] The entire list of sugary-dense product included in this calculation can be found in the appendix.

Imagine, for instance, that a household buys only sugary-dense products in these stores and purchases other food elsewhere. This household would then have a share of sugary-dense products on the total shopping basket close to 100%. To address this issue, the model focussing on observed heterogeneity in preferences is estimated on the subsample of households spending most of their food budget in this retail chain. That is, at least 80% of their food-at-home budget.[9] This leaves 130,514 households and around 4.2 million shopping trips where they bought a cola. Table 2.1 below shows the summary statistics about household demographics in the entire sample. While Table 2.2 displays some information about the share of sugary-dense products and shopping trips in the subsample of frequent costumers.[10]

Table 2.1: Household demographics (entire sample)

Variable	*Mean*	*Std. Dev.*	*Min.*	*Max.*
Income	24,447.72	4,098.37	5,836	82,575
Education	0.26	0.18	0	0.85
Household members	3.23	1.28	1	9
Kids	0.53	0.50	0	1
Age > 70	0.07	0.25	0	1
70 > Age > 58	0.14	0.35	0	1
58 > Age > 46	0.31	0.46	0	1
46 > Age > 34	0.34	0.47	0	1
Age < 34	0.13	0.34	0	1

Notes: *These variables are the following:* ***Income*** *- median income per capita of the district;* ***Education*** *- the percentage of higher educated people in the district;* ***Household members*** *-number of household members at the address level;* ***Kids*** *- dummy equal to 1 if there are children in the household, at the address level;* ***Age*** *- dummy by age group for the older member of the household, at the address level.*

Table 2.2: Sugar share and store visits (frequent costumers)

Variable	*Mean*	*Std. Dev.*	*Min.*	*Max.*
Sugar share	0.08	0.04	0.00	0.48
N° of store visits	32.38	30.64	5	300

Notes: *Frequent consumers are those spending at least the 80% of their food at home budget in these stores.*

The market considered is the one for bottled colas having at least 0.5% of market share. This market is composed of a set of seven brands made by three different producers. These brands

[9] The figures about the total food-at-home spending by household type are provided by Gfk Belgium.

[10] Figure 2.A.1. (in the appendix) shows the distribution of the share of sugary products across households.

can be sold in different bottle sizes: 1 litre, 1.5 litres, or 2 litres. Just two of them (brand A and B) are present in all these three different formats. In total, there are fourteen different cola products. Two of these can be considered as discount products (Brand D.1 and D.2), while others may be perceived as high-quality brands. Six of these colas contain sugar while the others do not (diet colas). Sometimes the retailer or the producer applies quantity rebates where the unit price of a product decreases conditionally on buying a certain quantity. The dataset also includes information about which product was included in the weekly folder of suggestions and deals that is sent home to customers. There is full information about which product was on the supermarket shelves at a given day and which one was out of stock. This information allows controlling for choice set heterogeneity across stores over time. Table 2.3 below summarizes the characteristics of the products analyzed. The unit of observation here is a household shopping trip. Hence, market shares are computed as the percentage of shopping trips where households bought a given brand.

Table 2.3: Product characteristics

Product	*Sugar*	*Size (L)*	*Not in store*	*Price/L (μ)*	*Price/L (Min)*	*Price/L (Max)*	*Market Share*
Brand A.1	sugary	1	2.88%	1.57	0.92	1.70	12.91%
Brand A.2	sugary	1.5	2.52%	1.18	0.57	1.39	18.46%
Brand A.3	sugary	2	51.63%	1.08	0.65	1.39	5.78%
Brand B.1	diet	1	10.74%	1.55	0.77	1.70	9.86%
Brand B.2	diet	1.5	3.47%	1.18	0.57	1.39	17.97%
Brand B.3	diet	2	65.84%	1.08	0.65	1.37	2.45%
Brand C.1	diet	1	25.08%	1.59	0.92	1.70	2.75%
Brand C.2	diet	1.5	4.96%	1.22	0.57	1.39	6.58%
Brand D.1	sugary	1.5	4.11%	0.32	0.23	0.40	5.32%
Brand D.2	diet	1.5	4.37%	0.30	0.17	0.39	2.45%
Brand E.1	diet	1	69.79%	1.26	0.89	1.35	0.74%
Brand E.2	diet	2	3.49%	0.76	0.42	1.05	12.43%
Brand F.1	sugary	1.5	44.33%	1.24	0.68	1.36	0.99%
Brand G.1	sugary	1.5	50.90%	0.77	0.45	0.98	1.31%

Notes: *The third column shows the package size of each item in terms of litres of product in a bottle. The fourth column, "Not in-store," shows the percentage of observations where the product was not available in the store. The fifth, sixth, and seventh column display the mean, the maximum and the minimum price per litre respectively.*

Unlike the conventional scanner price data used in the literature (e.g., Nielsen measured prices), this study uses more detailed data on posted prices in each store. The advantage of using posted prices is that they are not conditional on purchase and thus less sensitive to local and cyclical shocks (Coibion et al., 2015). They are not dependent neither on measurement errors due to loyalty cards (Einav et al., 2010). The retailer automatically collects posted prices daily for every item sold in each store of the group. In this way, it is possible to recover the full set of cola prices faced by the household during each shopping trip. These prices vary considerably both across stores and overtime. This is because these retailers are publicly committed to act as local price followers. Hence, their prices are matching those of other local competing supermarkets.[11] Figure 2.A.2 and 2.A.3 in the appendix show the price dispersion over stores and time, respectively.

2.4 The Cola Choice Model

Consider a model where each household i, at any shopping occasion $t = 1,2, \dots, T_i$, chooses one cola j from a time-varying choice set of J_t different colas. Each household maximises its utility by choosing the cola that delivers the highest utility of consumption. Let u_{ijt} represents the utility of household i conditional on buying product j on purchase occasion t and $\bar{U}$ denote the highest utility over all the alternatives in the set J_t. Then, the household i at time t will purchase good j only if:

$$u_{ijt} > \bar{U}.$$

The conditional utility u_{ijt} can be decomposed into $u_{ijt} = V_{ijt} + \varepsilon_{ijt}$. Where V_{ijt} represents the part of utility detectable from the observable data (to the researcher) and ε_{ijt} is the random error term, which is assumed be iid according to a type I extreme-value distribution (McFadden, 1974).[12] Let the conditional utility u_{ijt} be given by:

$$u_{ijt} = \alpha_{j_B} + \pi_i(Size)_j + \delta_i(Sugary)_j + \theta_i(Rare)_j - \beta_i(Price)_{ijt} +$$

$$+\varphi_i(Rebate)_{jt} + \omega_i(Folder)_{jt} + \gamma_i I(d_{t-1})_{ijt} + \varepsilon_{ijt}. \quad (2.1)$$

The first four terms of the utility specification in 4.1 are all those characteristics that are product specific and time-invariant. α_{j_B} is the brand-specific constant capturing all those factors that are brand specific and that do not change over time. These may be generally interpreted as the

[11] Hindriks & Serse (2019) use spirit price data from the same supermarket chain and show a significant price dispersion over space due to the price-matching strategy of this group.

[12] This is a typical assumption for Logit models, and it is maintained for the rest of the paper. Although this assumption seems very restrictive as it rules out correlation of unobservable factors across products, it is quite convenient for the resulting closed-formed choice probabilities. Correlation in unobserved factors across products is however introduced in the Mixed Logit version of the model through random coefficients for households' preferences in product attributes.

brand image or appreciation. $Size_j$ is the bottle's size measured in litres. Its coefficient π_i captures the household preference for a larger bottle size. $Sugary_j$ is a dummy variable equal to one if cola j contains sugar. The coefficient δ_i measures the household taste for sugary colas, which can be either negative (it dislikes sugar) or positive (it likes sugar). As not every cola is always available, the model includes the term $Rare_j$, which is a dummy variable indicating whether cola j is rarely present in the supermarket shelves (i.e., missing more than one-third of the time). Its coefficient θ_i captures the household appreciation for less commonly available colas once they are purchasable in the store.

All those time-varying variables affecting the household utility u_{ijt} are placed from the fifth to the eighth term of Eq. 2.1. $Price_{ijt}$ is the posted price of cola j faced by household i during the shopping trip t.[13] Its coefficient β_i measures the price sensitivity of household i, which is assumed to be positive for every households (i.e., $\beta_i \geq 0, \forall i$). This means that every household either dislike or is indifferent to higher prices. $Rebate_{jt}$ and $Folder_{jt}$ are dummy variables equal to one if cola j was on quantity rebate or in the weekly folder during the shopping trip t.[14] Their associated coefficients, φ_i and ω_i, measure the household sensitivity to rebates and the inclusion in the weekly folder, respectively. The term $I(d_{t-1})_{ijt}$ is an indicator variable that equals one if household i purchased cola j during its previous shopping trip (i.e., at time $t-1$). Its coefficient γ_i captures the household state dependence in cola choice. If $\gamma_i > 0$, the household exhibits positive state dependence. That is, the probability of choosing cola j increases when cola j is purchased in the previous shopping trip. Conversely, if $\gamma_i < 0$, the household is variety-seeker. It likes to try different colas, so the probability of choosing cola j decreases if cola j was purchased in the previous shopping trip.

The model does not include any outside option. This allows focusing on the purchase probabilities of each brand conditional on buying a cola, thus circumventing the potential issues arising from cola stockpiling (see Hendel & Nevo, 2006; Wang, 2015). Imposing an outside good without accounting for stockpiling dynamics can lead to an incorrect estimation of the intensive margin of demand (i.e., how much colas are bought). Since the focus of the paper is to investigate state dependence and taste heterogeneity in cola choice, avoiding more complex modelling and not focusing on the intensive margin can facilitate the identification of these two

[13] Prices are indexed by household since they vary across stores. Thus, two households can face different prices if they shop during same day but in two different stores.

[14] These two variables are not indexed by household since they do not vary over stores.

behaviors. Furthermore, as taxes would target only sugary varieties, the actual outside goods when performing tax simulations are diet colas.

One implication of this approach is that it does not permit variations in the overall cola expenditure overtime, but it permits variations in the share of sugary colas purchased. Therefore, the price elasticity of demand indicates how price changes lead households to substitute from one cola to another regardless of how much cola they buy. This is still appropriate for analyzing the corrective impact of sugar taxes, since leading households to switch from a sugary to a diet cola can account for a significant reduction in sugar intake. The main assumption here is that once a tax is implemented most of the substitution to other products occurs within the cola segment. That is, consumers tend to substitute sugary colas with diet colas, which are essentially the same product but without sugar content. Nevertheless, Table 2.A.7 in the appendix shows the results of a model where the purchase of another soda is considered as outside option. Interestingly, the demand parameters do not vary much from the baseline model with no outside good, suggesting that the introduction of an outside option does not invalidate the model results.

2.4.1 Homogeneous Tastes

Depending on the assumptions about the distribution of these taste coefficients, we can specify two different types of MNL models: 1) the simple Logit and 2) the Mixed Logit. Let the vector of household i's taste coefficients be denoted by $\vartheta_i = \left(\alpha_{j_B},\ \pi_i,\ \delta_i,\ \theta_i,\ \beta_i,\ \varphi_i,\ \omega_i,\ \gamma_i\right)$. In the simple Logit model, taste coefficients are assumed to be fixed across households. Hence, the index i in the vector of taste coefficients ϑ_i can be dropped, thus becoming $\vartheta = \left(\alpha_{j_B},\ \bar{\pi},\ \bar{\delta},\ \bar{\theta},\ \bar{\beta},\ \bar{\varphi}\ ,\ \bar{\omega},\ \bar{\gamma}\right)$. Assuming the error term ε_{ijt} to be iid according to a type I extreme-value distribution, the probability that household i chooses good j at time t is then given by the following Logit formula (McFadden, 1974):

$$L_{ijt} = \frac{e^{\vartheta' X_{njt}}}{\sum_j e^{\vartheta' X_{njt}}}. \tag{2.2}$$

Where X_{njt} is the vector of product characteristics and the other explanatory variables in Eq. 2.1, which multiplies the vector of taste coefficients ϑ. The aggregated market share of good j, s_j, is the average of the choice probability L_{ijt} across each household i and shopping trip t:

$$s_j = \frac{\sum_i \sum_t L_{ijt}}{N}.$$

Where N is the total number of observations.[15] Thus, by assuming the set $J = \{j_1, j_2, \dots j_n\}$ to cover the entire market of interest, we would have that $\sum_j s_j$=1. Own and cross-price elasticity of the Logit demand for good j can be derived by taking the derivative of s_j with respect to the price of cola z (denoted by $Price_z$):

$$\eta_{jz} = \frac{\partial s_j}{\partial (Price)_z} \frac{(Price)_z}{s_j} = \begin{cases} -\beta (Price)_j (1 - s_j) & if\ j = z \\ \beta (Price)_z s_z & if\ j \neq z \end{cases} \tag{2.3}$$

Although this model has the advantage of having a simple closed-form solution for the choice probabilities, it imposes very restrictive and unrealistic substitution patterns. This is because the relative choice probability for any couple of goods does not change for any variation in attributes of a third product (i.e., IIA property). Another limitation is that this model does not account for the correlation of errors for each household over time. Which is likely to arise in case of repeated purchases (i.e., panel data). The presence of dynamics in unobservable factors can induce inconsistency in the estimation of state dependence, as the lagged dependent variable is correlated with the current error term (Train, 2009). Foremost, the assumption of homogeneous taste across households can heavily restrict the demand analysis since it only focuses on the average impact of changes in product attributes on demand.

Although the restrictive substitution patterns and homogeneity in taste parameters, the Logit model has the advantage of being easy to estimate on large samples. Hence, this model can be estimated on the entire dataset, which includes 434,475 households and around 9.7 million shopping trips. The results of this estimation can give a first insight into the average household taste, price sensitivity, and state dependence by including as many observations as possible. The model is estimated through Maximum Likelihood procedure, with the likelihood function taking the following form (McFadden, 1974; Train, 2009):

$$LL(\vartheta) = \prod_i \prod_t \prod_j (L_{ijt})^{y_{ijt}}.$$

Where y_{ijt}=1 if household i chooses good j at time t and 0 otherwise. ϑ is the vector of fixed taste parameters to be estimated. Its estimator is the gradient vector that maximizes $LL(\vartheta)$.

2.4.2 Unobserved Taste Heterogeneity

While the simple Logit assumes that taste coefficients are homogeneous across households, the Mixed Logit model allows them to vary randomly in the population. This model has the

[15] One observation is intended as a single shopping trip made by a household. Thus, N is equal to the sum of shopping trips T_i over all households in the sample.

advantage of accounting for the panel structure of the dataset by allowing unobserved factors that are household-specific to be correlated over different shopping trip *t*. In the Mixed Logit, the probability that household i makes a given sequence of choices, conditional on its vector of taste coefficients ϑ_i, is given by the following product of logit formulas (McFadden, 1974; Train, 2009):

$$L_{ij}(\vartheta_i) = \prod_{t=1}^{T_i} \left(\frac{e^{\vartheta_i' X_{njt}}}{\sum_j e^{\vartheta_i' X_{njt}}} \right). \tag{2.4}$$

Where X_{njt} is the vector of product characteristics and the other explanatory variables that multiplies the vector of household-specific taste coefficients ϑ_i. Contrary to the simple Logit, this specification allows price sensitivity β_i, taste for sugar δ_i, and preference for package size π_i to vary randomly in the population and to be correlated with each other. In particular, β_i is assumed to be distributed as normal truncated at zero (i.e. $\beta_i \sim TN_{[0,+\infty]}[\bar{\beta}, \sigma_\beta^2]$) so that no households like higher prices. While both δ_i and π_i are assumed to be normally distributed (i.e. $\delta_i \sim N[\bar{\delta}, \sigma_\delta^2]$ and $\pi_i \sim N[\bar{\pi}, \sigma_\pi^2]$), which means that they either like or dislike sugary colas and larger bottle size. A standard assumption in Mixed Logit model is that the random taste coefficients are independent of product attributes. That is, ϑ_i is independent of X_{njt}. This assumption implies that households' preferences are exogenous (i.e., they cannot be endogenously determined by changes in product characteristics). The other taste coefficients are assumed to be fixed to ease the estimation procedure. Choosing these three coefficients to be randomly distributed allows for very flexible substitution patterns, as colas that are similar in price, sugar, size would be closer substitutes. As a result, the vector of household i's taste coefficients can be now denoted as $\vartheta_i = \left(\alpha_{j_B}, \pi_i, \delta_i, \bar{\theta}, \beta_i, \bar{\varphi}, \bar{\omega}, \bar{\gamma}\right)$. Since ϑ_i is not observable, the unconditional choice probability will be the integral of $L_{ij}(\vartheta_i)$ over all possible values of ϑ_i:

$$s_j = \int L_{ij}(\vartheta)\, f(\vartheta|W)\, d\vartheta. \tag{2.5}$$

Where $f(\vartheta|W)$ is the joint density functions of the taste coefficients $(\beta, \delta, \pi)^{\intercal}$. With ϑ and W being the parameters of these distributions. These taste coefficients can be also allowed to be correlated with each other.[16] s_j is the average probability of choosing good j and can be

[16] In particular, correlation in taste coefficients is allowed in the basic specification without observed heterogeneity.

interpreted as the market share of the alternative j. From Eq. 2.5 we can derive the price elasticity of demand for the alternative j with respect to any other alternative z:

$$\eta_{jz} = \frac{\partial s_j}{\partial (Price)_z} \frac{(Price)_z}{s_j} = \begin{cases} -\frac{(Price)_j}{s_j} \int \beta_i L_{ij}(\vartheta)\left[1 - L_{ij}(\vartheta)\right] f(\vartheta|W)\, d\vartheta\ , & if\ j = z \\ \frac{(Price)_z}{s_j} \int \beta_i L_{ij}(\vartheta) L_{iz}(\vartheta)\, f(\vartheta|W)\, d\vartheta\, , & if\ j \neq z. \end{cases} \tag{2.6}$$

The upper term in the bracket represents the own-price elasticity of demand (i.e., the percentage change in the probability of purchasing alternative j given a percentage change in its price). While the lower term in the bracket is the cross-price elasticity of j with respect to z (i.e., the percentage change in the probability of purchasing alternative j given a percentage change in the price of alternative z). The product shares do not necessarily drive these price elasticities as in Eq. 2.3. Here, every household have different price sensitivity β_i. The price elasticity of the aggregate demand can therefore be interpreted as the mean price sensitivity across households, weighted by their specific purchase probabilities. The correlation between $L_{ij}(\vartheta)$ and $L_{iz}(\vartheta)$ drives the substitution patterns over different values of ϑ. This means that products having similar characteristics, such as price, sugar content, and package size, will be closer substitutes to each other.

This approach can provide precious insights not only on the average “treatment effect” of a given policy but also on its heterogeneous impact across households even in the absence of demographic data. As the model can capture heterogeneity in price sensitivity across households, it also allows estimating the curvature of aggregate demand. Which is a fundamental variable in predicting tax pass-through in imperfectly competitive markets (Weyl & Fabinger, 2013). Contrary to the Logit, this specification accounts for correlation in unobserved factors over time and can provide a robust estimate of state dependence coefficient $\bar{\gamma}$ as it controls for permanent taste heterogeneity across households.[17]

A limitation of the Mixed Logit, however, is its heavy computational burden as compared to the simple Logit. The estimation procedure often requires a large amount of RAM, which makes such models challenging to estimate on vast datasets. For the ease of computation, this model is estimated on a random subsample of 10,000 households through Simulated Maximum Likelihood. Following Train (2009), the estimation procedure can be summarized as follows.

[17] Accounting for correlation in the household unobservable factors overtime through random taste coefficients does not necessarily rule out the possibility of autocorrelation in the remaining error terms (see, for instance, Keane, 1997; Dubé et al., 2010). This issue is addressed in section 7.

First, it is taken a draw of ϑ from $f(\vartheta|W)$ and it is the computed the following logit formula $e^{(\vartheta_i' X_{njt})} / \sum_j e^{(\vartheta_i' X_{njt})}$ for each shopping trip t. Second, it is taken the product of these logit formulas in order to get the simulated probability $L_{ij}(\vartheta^R)$ with this first draw. Third, these two steps are repeated for R times and results are averaged. This exercise gives the following average simulated probability:

$$\tilde{L}_{ij} = \frac{1}{R}\sum_{r=1}^{R} L_{ij}(\vartheta^R). \tag{2.7}$$

Which is then plugged into the log-likelihood function to obtain the simulated log-likelihood (SLL):

$$SLL = \sum_{i=1}^{I} \ln\left\{\prod_{j=1}^{J} \tilde{L}_{ij}{}^{y_{ijt}}\right\}. \tag{2.8}$$

Where $y_{ijt} = 1$ if the household i purchase good j in shopping occasion t and zero otherwise. The SLL function is then maximized to obtain the parameters W of the density functions $f(\vartheta|W)$ of the taste coefficients.

2.4.3 Identification

A key parameter to be estimated is the price sensitivity coefficient β_i. Which measures the causal impact of price changes on the probability of choosing a cola product. Similarly to Dubois et al. (2020), this parameter is estimated by exploiting two sources of price variation. The first is the cross-store variation in prices and choice set over space. The second is the price variation within the brand across different bottle sizes, which can vary over both stores and time. These two sources of price variation deliver a substantial heterogeneity in price vectors and choice sets faced by households over different shopping trips. The price vector and choice set considered for each household are those of the store they visited on a given shopping trip. As mentioned in the data section, these two vary over both stores and time daily. Hence, there is complete information about prices and choice set faced by the household during that day.

The fact that each store is committed to match the lowest price in its local market ensures enough variability in prices over choice occasions. Moreover, this price-matching strategy reduces the problems arising from households selecting the store based on the price level of their favourite brand across different chains. These households could be problematic because they would appear in the sample only when the price of their favourite cola is the lowest in this retail chain. Hence, making their choices based on prices of other retailers that are unobserved.

The local price-matching strategy limits this issue because the observed vector of prices in these stores reflects those of other retail chains in the area.

The primary identifying assumption is that the set of explanatory variables X_{ijt} is not correlated with the random shock ε_{ijt}. In particular, a correct identification of the price coefficient β_i requires that systematic demand shocks for different brand types do not drive cola price changes. In other words, price variations must be exogenous to household demand. This assumption seems quite reasonable in this setting, given the daily frequency of price updates and the household level demand data. However, this assumption would be invalid if stores can anticipate brand-specific demand shocks and change prices accordingly. In this study, such a scenario is presumed to be unlikely. Instead, it is assumed that variations in prices are the resultant of different cost shocks and random temporary price discounts that are induced by the retailer price-matching strategy. To give an idea of such price variation, Figure 2.1 displays the price evolution of one brand in a given store over time.[18] As shown in the figure, the price series displays a large number of temporary price cuts.[19] The assumption here is that these price cuts are not the resultants of stores anticipating temporary negative taste shocks for a given brand. They are instead assumed to occur due to random temporary price promotions induced by the retailer's price-matching strategy.

Figure 2.1: Example of price variation

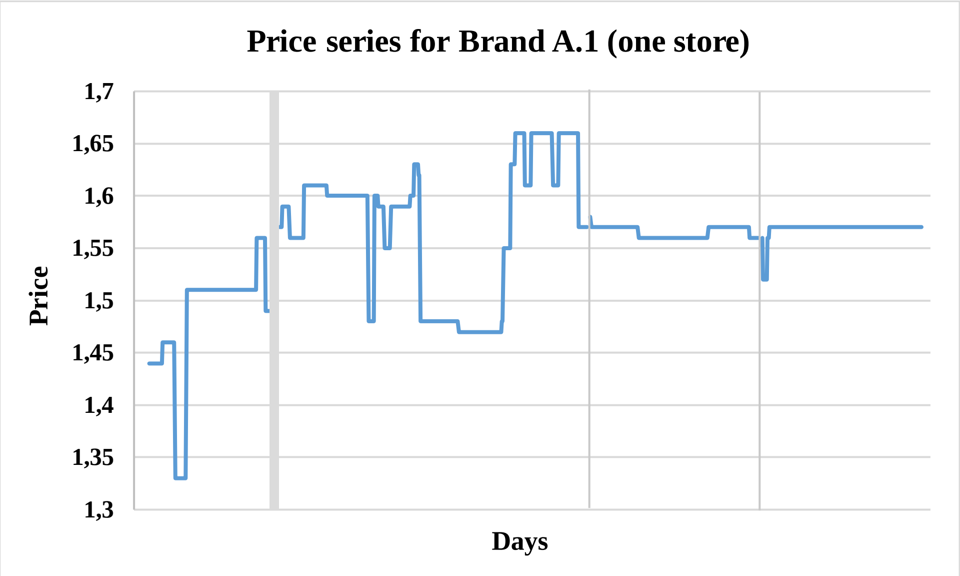

[18] The grey areas in the graphs indicate that the product was not available (not on shelves) during those days in that store.

[19] Each brand-store specific price series has a different shape. This is because price discounts are not necessarily coordinated across stores as each establishment follows the prices in its local market.

Another critical assumption for the identification of the price coefficient is that price changes are not correlated with unobserved product characteristics. The extensive set of control variables such as rebates, folder information, and product availability limits this issue as they control for some characteristics of the store environment that are typically unobserved with other types of scanner data. Yet, it cannot be ruled out that omitted variables in the utility specification are correlated with prices. In a further specification, brand-specific effects are allowed to vary monthly to control for possible temporary unobserved shocks at the brand level that are common across all stores. This approach controls notably for adverting campaigns, which are mostly done nationally and are not observed in the data.[20]

In addition to price sensitivity, this work focuses on identifying household taste heterogeneity and state dependence in product choice. The identifying assumption for taste heterogeneity is that persistent household preferences for colas can be decomposed into a vector of preferences for three product characteristics. Which are price, sugar, and bottle size. These preferences are assumed to be parametrically distributed in the population and can vary with observed household characteristics. State dependence is identified as systematic deviations from the persistent household preferences that are driven by the previous shopping experience (lagged choice variable). The variation in prices and choice set overtime is fundamental in disentangling taste heterogeneity from state dependence. These temporary market shocks can in fact lead the household to switch from its favourite product to another cola. In the absence of state dependence, the household should revert to its favourite cola once these shocks are over. However, if the household keeps on choosing the last product purchased, then there is evidence of dynamics in cola choice that is driven by the recent shopping history of the household.

The correct identification of state dependence requires ruling out other possible sources of choice inertia that might be confounded with a real structural dynamic in product choice (Heckman, 1978). If the underlying distribution of taste heterogeneity across the population is not correctly specified, there is a risk of capturing some spurious state dependence that is due to some persistent unobserved propensity of the household in making a given choice. Moreover, even if the household heterogeneity in taste is correctly specified, it is still possible that the estimated choice dynamics are due to autocorrelated taste shocks. In this scenario, there can be

[20] Time-varying brand effects are not included in the main specification as they can increase the computational burden, hence making the tax simulation exercise more cumbersome. However, as shown in Section 2.7, the price coefficient does not change much when including these controls.

a spurious state dependence just because the lagged choice variable is capturing a large random utility draw on both the previous and current shopping trip.

As denoted by Dube et al. (2010), the implications of this model are remarkably different from those of a model with structural state dependence. In a model with autocorrelated errors, choice dynamics are generated by unobserved factors and hence cannot be influenced by any marketing variable. In the context of sugar taxes, this means that governments cannot exploit state dependence through a tax hike to generate a larger reduction in cola demand over time. The issues of misspecified heterogeneity and autocorrelated error terms are addressed in Section 2.7, where various robustness checks are performed in order to rule out these alternative explanations for the observed choice dynamics.

2.5 Results

2.5.1 Demand Estimates

Table 4 shows the parameter estimates for both the Logit and Mixed Logit version of the utility specification in Eq. 2.1. Each of them is estimated with and without the state dependence variable (lagged choice). The models without the state dependence variable are labelled as *static* models. Their estimates are displayed in columns (1) and (3). While those accounting for state dependence are labelled as *dynamic* models, and their estimates are shown in columns (2) and (4). The first two columns of Table 2.4 show the parameter estimates of the Logit model, which assumes taste homogeneity across households. The last two columns display the parameter estimates of the Mixed Logit model, which assumes heterogeneity in taste across households instead. As the Mixed Logit model allows for price, sugar, and size parameters to be randomly distributed, these columns also include estimates of their standard deviations. Which are labelled σ_β, σ_δ and σ_π respectively. The coefficients $\bar{\beta}$, $\bar{\delta}$ and $\bar{\pi}$ are the estimates of their means. All models are estimated on a random subsample of 10,000 households. The Logit models are also estimated on the entire dataset and results are compared to those obtained with the random subsample. These results can be found in table 2.A.1 in the appendix. The comparison of the parameter estimates for the Logit models across different samples indicates very similar results. Thus suggesting that the Logit results are not sensitive to the reduction in sample size.

Firstly, I discuss the results under the assumption of taste homogeneity. Column (1) and (2) show a positive price coefficient $\bar{\beta}$ ranging from 0.82 to 0.45. Which indicates that households dislike higher prices.[21] Taste for sugar $\bar{\delta}$ and bottle size $\bar{\pi}$ are also positive and below one. This means that the average household prefers sugary to diet colas and likes larger bottles. The estimate for the rarity coefficient $\bar{\theta}$ instead indicates that households dislike colas that are rarely on shelves, once they are available. This may suggest that a constant exposure on supermarket shelves tends to increase product appreciation. The coefficient for folder $\bar{\omega}$ is slightly negative in the *static* specification, while it is slightly positive in the *dynamic* one. The rebate coefficient $\bar{\varphi}$ is positive in both specifications but close to zero. This indicates that both quantity rebates and promotions in the weekly customer folder increase the demand but only to a small extent.

Table 2.4: Parameter Estimates

	Taste Homogeneity		**Taste Heterogeneity**	
Parameter	***Static* (1)**	***Dynamic* (2)**	***Static* (3)**	***Dynamic* (4)**
Price $\bar{\beta}$	**0.816** (0.028)	**0.450** (0.037)	**4.441** (0.040)	**2.531** (0.053)
Price σ_β	-	-	**6.118** (0.031)	**4.132** (0.040)
Sugary $\bar{\delta}$	**0.782** (0.016)	**0.492** (0.018)	**1.266** (0.027)	**0.782** (0.030)
Sugary σ_δ	-	-	**3.818** (0.018)	**2.428** (0.019)
Size $\bar{\pi}$	**0.291** (0.021)	**0.175** (0.026)	**-0.459** (0.022)	**-0.441** (0.024)
Size σ_π	-	-	**5.441** (0.018)	**2.977** (0.019)
Rare $\bar{\theta}$	**-1.039** (0.011)	**-0.500** (0.013)	**-1.958** (0.012)	**-0.973** (0.013)
Folder $\bar{\omega}$	**-0.014** (0.007)	**0.046** (0.010)	**-0.093** (0.010)	**-0.026** (0.011)
Rebate $\bar{\varphi}$	**0.079** (0.007)	**0.064** (0.009)	**0.176** (0.009)	**0.163** (0.010)
State Dependence $\bar{\gamma}$	-	**3.033** (0.005)	-	**1.854** (0.005)
Model	Logit	Logit	Mixed Logit	Mixed Logit
McFadden's R^2	0.09	0.48	0.46	0.55

Notes: *standard errors are in parenthesis. All parameters are significant at the 1% level.*

[21] In Eq. 2.1 price enters the utility as a negative term. Hence, a positive coefficient means a lower utility if price increases.

Interestingly, the state dependence coefficient $\bar{\gamma}$ is strongly positive. This finding suggests that having purchased a cola in the previous shopping trip increases the probability of repurchasing the same cola on the next shopping occasion. Furthermore, the lagged choice variable is a significant predictor of the model. The inclusion of the lagged choice variable, in fact, increases the pseudo R^2 from 0,09 to 0,48. Hence, showing that under the assumption of homogeneous tastes, the *dynamic* logit specification is superior to the *static* one.

Accounting for unobserved taste heterogeneity, however, delivers a complete picture of households' preferences. As shown in column (3) and (4), all standard deviations of the randomly distributed taste coefficients are large and significant. This result rejects the Logit assumption of homogeneous taste coefficients, indicating that households do have heterogeneous preferences for colas. The *dynamic* Mixed Logit model fits the data better than any other specification, with a pseudo R^2 equal to 0.55. This is also confirmed by a likelihood test ratio between this model and each of the other three specifications. Thus, suggesting that households exhibit both state dependence and taste heterogeneity. The comparison between column (3) and (4) shows that the omission of the state dependence variable in the Mixed Logit specification leads to an overestimation of the unobserved taste heterogeneity across households. The standard deviations of the randomly distributed taste coefficients in column (3) are indeed larger (about one third) than those displayed in column (4). Moreover, price sensitivity decreases considerably once accounting for state dependence. Hence, ignoring such behavior would overestimate the direct impact of prices on demand.

The *dynamic* model with unobserved taste heterogeneity is estimated using 200 Halton sequences. Increasing the number of Halton sequences delivers stable parameter estimates. The results of these computations are shown in table 2.A.2 in the appendix. The price parameter of this model, which was assumed to be distributed as truncated (at zero) normal, has a mean of 2.53 and a median equal to 1.53.[22] The estimate of its standard deviation σ_β indicates that households are quite heterogeneous in terms of price sensitivity. The 12% of households have a price coefficient below 1 and, therefore, can be considered as relatively insensitive to prices. On the other hand, 13% percent of households have a price coefficient higher than 8, and thus can be considered as extremely sensitive to prices. This indicates that households are, on average, more price-sensitive compared to the estimates of the Logit specification. Hence,

[22] The model was also estimated assuming either normal or lognormal distribution for the price coefficient. Yet, in the first case, a quite large part of household had a negative price parameter. While in case of log-normality, the estimates were not stable when increasing the number of random draws.

assuming the price coefficient to be homogeneous across households leads to an underestimation of the average price sensitivity.

The mean of the sugary parameter $\bar{\delta}$ is quite similar to the Logit specification, confirming that on average households like sugary colas. Yet, the estimate of its standard deviation σ_{δ} also highlights that taste for sugar is quite heterogeneous in the population. Similar results are found in preferences for bottle size, which are also quite heterogeneous across households. Most people prefer small bottles, but a large part likes colas in larger formats. Interestingly, the state dependence coefficient $\bar{\gamma}$ displayed in column (4) drops considerably compared to the one displayed in column (2). However, this remains positive and significant, decreasing from 3.03 to 1.85. This is because the Mixed Logit controls for persistent taste heterogeneity. While in the Logit specification, this can be empirically confounded with a spurious state dependence. Therefore, the Mixed Logit estimate of $\bar{\gamma}$ provides more robust evidence of state dependence in cola choice. Suggesting that part of the state dependence found with the Logit specification was instead due to unobserved taste heterogeneity across households.

2.5.2 Demand curves

The parameter estimates allow computing the impact of price changes on cola demand. Table 2.5 displays the matrix of own and cross-price elasticities of demand using the *dynamic* Mixed Logit model (computed as in Eq. 2.6). As the model does not include any outside option, these elasticities indicate the percentage change in the average purchase probability conditional on purchasing a cola for a 1% change in price. In other words, they provide the measure in which prices lead households to substitute from one cola to another, regardless of how much cola they buy. Interestingly, the impact of a change in price on the average purchase probability tends to increase in the long-run due to state dependence in cola choice. Households need time to adjust their demand to a price change and reach the new equilibrium in their purchase probabilities.

The average price elasticity of demand across all colas after the first price iteration is equal to 1.74.[23] This is computed by simulating for each product the impact of a 1% increase in its price and keeping the household past purchase probabilities fixed at their level before the price increase. This measure can be interpreted as the short-run price elasticity of demand. Which considers the sole initial price effect and ignores the dynamics that will arise due to a change in the state dependence variable. Once the new purchase probabilities under the new price vector are reiterated in place of the old ones, household demand will continue to adjust until reaching

[23] This price elasticity is computed as a weighted average of all colas by using the average purchase probability as weight.

the new equilibrium shares. At the end of this process, the long-run (or final) average price elasticity of demand across all products is equal to 2.04. This elasticity is larger than the short-run price elasticity since an initial price increase leads to a reduction in the current choice probability, which in turn reduces the purchase probability in the next shopping occasion.

Table 2.5: Own and cross-price elasticities of demand (*Dynamic* Mixed Logit)

	A.1	A.2	A.3	B.1	B.2	B.3	C.1	C.2	D.1	D.2	E.1	E.2	F.1	G.1
A.1	**-1.35**	0.06	0.25	0.59	0.10	0.05	0.01	0.02	0.02	0.05	0.09	0.03	0.02	0.07
A.2	0.32	**-1.59**	0.71	0.08	0.25	0.08	0.01	0.16	0.02	0.02	0.28	0.07	0.07	0.21
A.3	0.26	0.14	**-2.58**	0.07	0.11	0.03	0.00	0.04	0.01	0.01	0.06	0.02	0.01	0.06
B.1	0.30	0.02	0.04	**-1.59**	0.13	0.03	0.01	0.03	0.03	0.08	0.02	0.04	0.04	0.02
B.2	0.36	0.17	0.38	0.90	**-2.13**	0.03	0.04	0.40	0.07	0.12	0.09	0.24	0.31	0.08
B.3	0.46	0.14	0.16	0.15	0.06	**-1.87**	0.27	0.04	0.33	0.42	0.63	0.16	0.18	0.43
C.1	0.05	0.02	0.02	0.14	0.06	0.11	**-2.02**	0.04	0.26	0.34	0.02	0.13	0.15	0.04
C.2	0.08	0.12	0.15	0.19	0.42	0.03	0.04	**-1.63**	0.05	0.06	0.03	0.14	0.16	0.05
D.1	0.14	0.03	0.04	0.40	0.17	0.28	0.61	0.11	**-1.99**	0.97	0.12	0.38	0.46	0.11
D.2	0.04	0.02	0.01	0.10	0.03	0.04	0.05	0.01	0.07	**-3.48**	0.01	0.04	0.03	0.02
E.1	1.34	0.72	1.06	0.42	0.41	0.72	0.16	0.22	0.21	0.30	**-2.44**	0.30	0.38	0.84
E.2	0.11	0.05	0.07	0.28	0.24	0.06	0.11	0.13	0.15	0.21	0.05	**-2.29**	0.23	0.06
F.1	0.32	0.14	0.20	0.91	0.72	0.15	0.35	0.40	0.44	0.63	0.24	0.59	**-2.45**	0.17
G.1	0.11	0.07	0.09	0.03	0.04	0.07	0.01	0.02	0.02	0.03	0.07	0.03	0.02	**-2.22**

Notes: *own-price elasticities are in bold characters. Elasticities are computed at the sample average across households.*

The Mixed Logit model allows price elasticities to depend on the distribution of taste and price sensitivity across households, with similar products being closer substitutes to each other. An easy way to capture this property is to look at the example displayed in Table 2.6. This table shows the percentage change in demand for different colas following a 10% increase in the price of Brand A.3, which is a bottle of sugary cola of 2 liters. As the price of this cola increases, households substitute for similar products in terms of sugar content and bottle size. A 10% increase in price for Brand A.3 leads to a reduction in its market share by 13.66%. Most of these consumers then switch other sugary and large colas (2 liters). Their market shares increase by 2.10% and 2.33%, respectively, while those for diet and small colas (1 liter) increase only to a smaller extent.

Table 2.6: Example of Substitution patterns for Brand A.3 (10% price increase)

	Brand A.3	*Sugary*	*Diet*	*Large (2L)*	*Small (1L)*
Δ ***market share***	-13.66%	+2.10%	+0.77%	+2.33%	+0.28%

Notes: *The table reports the change in the average purchase probability of sugary and diet colas following a 10% increase in the price of Brand A.3., which is a sugary cola of 2 liters.*

From the parameter estimates, it is possible to draw the demand curve for each cola in the market. Interestingly, as households have a different price elasticity of demand, demand curves are not bounded to be linear in price. This exercise can therefore give an approximation of the curvature of demand in terms of average purchase probability. The estimates of the demand curves for each product are shown in Section 2.A.3 of the appendix. For most products, demand curves are highly convex. As price increases, price elasticity of demand decreases. They are thus indicating that as price goes up, only the less price-sensitive households keep on purchasing the same product. The demand curves for the cheapest colas (D.1 and D.2), however, are much less convex than other brands. This is probably because households purchasing these colas are the most price sensitive. Therefore, they are much more likely to reduce their demand for higher price levels.

Importantly, these results cannot be directly extended to the shape of the aggregate demand for each product. Here the focus is on the average purchase probability conditional on buying a cola. It might be possible that for very high prices households either reduce the number of bottles purchased or decide not to purchase a cola. As a result, aggregate demand curves should be less convex than those displayed in Section 2.A.3. They might be similar only in the limit case where both the intensive margin and no purchase effects are equal to zero. The demand curves presented here can be interpreted as a lower bound in terms of price sensitivity of the aggregate demand for colas. They are thus providing the most conservative setting in order to evaluate the impact of price-based policies on the demand for sugary colas.

2.5.3 Tax Policy Simulations

The model estimates allow performing tax policy simulations to assess the impact of sugar taxes on the households' probability of purchasing sugary colas. Table 2.7 shows the outcome of different tax policy simulations on the average purchase probability of sugary and diet colas. The simulations assume full tax pass-through (i.e., 100%), and consider both ad-valorem and excise taxes. An ad-valorem tax on sugary colas is a tax levied as a percentage of their market price. Under this tax scheme, cheap colas will be taxed less in absolute value than more expensive brands. This type of taxation was recently implemented in various countries, although sometimes it affects all sodas regardless of their sugar content. For instance, in 2015, the Barbados implemented a 10% ad-valorem tax on SSB. While in Chile, since 2014, drinks with high sugar content are taxed by 18%. The main advantage of ad-valorem taxes is their ease of implementation, as the amount of tax levied varies only as a function of prices. However, their shortcoming is that they are not related to the amount of sugar contained in each product. Hence, they may be less effective both in reducing sugar intake among consumers and in incentivizing firms in reformulating their products. An alternative to this tax scheme is the implementation of excise taxes on sugar. These are taxes levied on the amount of sugar contained in each product so that prices increase proportionally to their sugar content. This type of tax was recently implemented in few American localities, such as Berkley, in 2015, and Philadelphia in 2017.[24] In the UK, since April 2018, sodas with more than 8 grams of sugar per 100 millilitres of products are taxed by £0.24 per litre. Another interesting example is Ecuador, where on SSB with less than 25 grams of sugar per litre is levied an ad-valorem tax of 10%. However, SSB with higher sugar content are instead taxed by an equivalent of $0.18 for every 100 grams of sugar.

Table 2.7: Sugar tax simulations

		Ad-valorem tax		***Excise tax***	
Δ *market share*	**Timing**	**20%**	**30%**	**€0,20 (/100gr)**	**€0,30 (/100gr)**
Sugary colas	**Δshort run**	-11.70% (€1.34 M)	-16.76% (€2.01 M)	-13.38% (€1.26 M)	-19.42% (€1.75 M)
	Δlong run	-13.13% (€1.32 M)	-18.69% (€1.97 M)	-14.47% (€1.24 M)	-21.29% (€1.70 M)
Diet colas	**Δshort run**	+11.37%	+16.28%	+13%	+18.87%
	Δlong run	+12.76%	+18.16%	+14.73%	+20.69%

Notes: *The table reports the change in the average purchase probability of sugary and diet colas under different sugar tax schemes. The simulations are performed at the sample average across households. Tax revenues are displayed in parenthesis.*

The results of the tax policy simulations displayed in Table 2.7 are expressed in terms of percentage change in the market share (i.e., average purchase probability) for the two different groups of cola. As sugary colas are the only type of product being taxed, their market share decreases considerably after any tax is implemented. The market share for diet colas instead

[24] In both Berkley and Philadelphia, the tax is not strictly levied on the amount of sugar, but it is a volumetric tax that is just levied on sugary beverages (i.e., diet sodas are not taxed).

increases, as households then substitute towards healthier varieties. Table 2.7 also shows the resulting government revenues for each tax rate.[25] These revenues are computed considering all shopping trips observed in the sample, which amount to roughly 9 million over a period of 2.5 years. In these simulations, tax revenues increase with the level of sugar taxation and tend to be larger when using ad-valorem taxes. Under perfect tax shifting, the 20% ad-valorem tax and the €0.20/100gr of sugar excise tax increase the average price of sugary colas similarly (from €1.48 to €1.78 and €1.80, respectively). The 30% ad-valorem tax and the €0.30/100gr of sugar excise tax do also increase the average price to a similar extent (from €1.48 to €1.93 and €1.95, respectively). Thus, for the same consumer price, excise taxes seem to be more effective in reducing sugary cola demand and raise less tax revenues.

Table 2.7 also shows the different impacts of these policies over time. Because of state dependence, the immediate impact of a price change on household demand tends to be smaller than in the long-run. Although the tax mostly affects demand already in the short-run, this finding suggests that the public health benefits of a sugar tax tend to increase over time. The difference between the short-run versus the long-run demand response to the tax, however, differs across tax schemes. This difference tends to increase with the initial tax change. The more the current purchase probability decreases due to a tax hike, the more the future purchase probability decreases due to a change in the state dependence variable.[26] Because of these dynamics, tax revenues tend to decrease in the long-run for the same tax rate. This difference is also increasing with the initial tax change. Thus, although tax revenues increase with the tax rate, in the long-run there would be a larger reduction in tax revenues for higher tax rates. This occurs as sugar taxes become increasingly more effective in the long-run due to state dependence in cola choice.

A typical concern with sugar taxes is that they may incentivise households to substitute for cheaper sugary brands. To study this possibility, Figures 2.2.a and 2.2.b below compare the effectiveness of ad-valorem versus excise taxes in reducing cola demand across different product types. The y-axes indicate the different levels of taxation, while the x-axes represent the market shares. The blue curve is the demand for the cheapest sugary cola in the market (brand D.1) as a function of the tax. The red curve is instead the demand for all other sugary

[25] As tax simulations are done using a conditional demand model, tax revenues here do not consider the intensive margin. They are, therefore, a lower bound for the generated tax revenues.

[26] As demand curves are quite convex, this difference is not strictly increasing in the tax rate. At a certain tax level, the difference between short-run and long-run demand response will not increase with a higher tax. This occurs when the tax reaches the inelastic part of the demand so that any further tax increase does not marginally affect market shares.

colas. Figure 2.2.a shows that the demand for the cheapest sugary cola is inelastic to ad-valorem taxes, while Figure 2.2.b shows that this is quite elastic to any level of excise tax. The demand for more expensive sugary colas instead declines similarly under both tax schemes.

Figure 2.2.a: ad-valorem tax

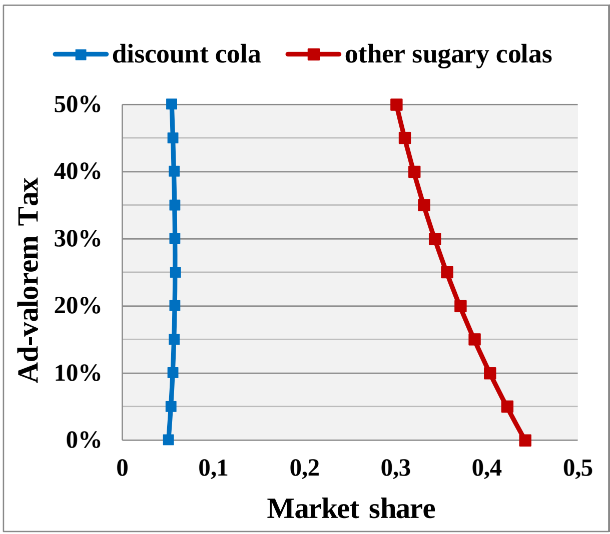

Figure 2.2.b: excise tax

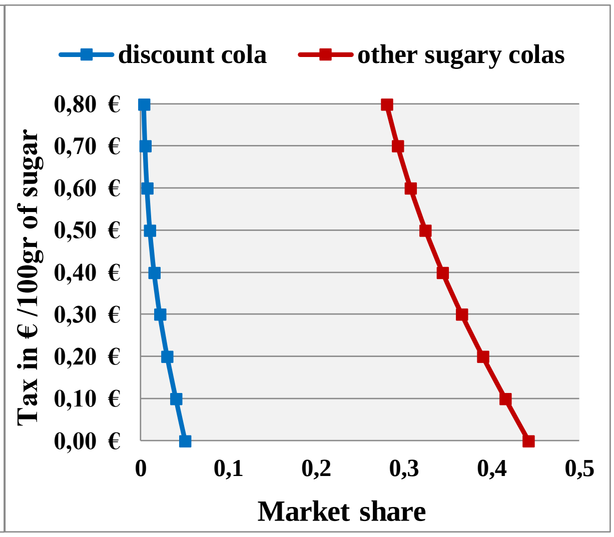

Notes: *Figures 2.2.a and 2.2.b show the impact of ad-valorem and excise taxes on the demand for sugary colas. The y-axes indicate the different levels of taxation. The x-axes indicate the market shares (i.e., average purchase probability). The blue curve represents the cheapest sugary cola in the market (brand D.1). The red curve represents all other sugary colas. The simulations are performed at the sample average across households. In these figures, only the long-run effect is considered.*

This difference in demand response is originated from the measure in which these two types of taxes affect the price of cheap brands. Ad-valorem taxes tend to increase the price of cheap colas by a smaller amount than specific taxes on sugar, thus leading to the following two effects. First, the least price-sensitive households purchasing the cheap brand are not concerned by such a small price increase and hence keep on purchasing this cheap product. Second, those households preferring expensive brands, having a strong taste for sugar, and high price-sensitivity are likely to substitute their preferred cola with another sugary cola but of a cheaper brand. The sum of these two effects makes the demand for the cheapest sugary cola inelastic to any level of ad-valorem tax, as shown in Figure 2.2.a above. When a tax is levied on sugar content, these two effects are reduced since the price increase for cheaper colas tends to be larger than with an ad-valorem tax. Furthermore, as all sugary colas have similar sugar content, a specific tax on sugar implies a similar price hike across sugary colas in absolute terms. This feature increases the relative price between cheaper and expensive sugary colas, hence disincentivizing the substitution towards cheaper brands. As a result, these findings suggest that specific taxes on sugar should be preferred to ad-valorem taxes to avoid a substitution towards cheaper sugary brands.

Nevertheless, as the focus is on conditional demand, the results of above should be taken with caution. A critical assumption is that the extensive margin reaction to ad-valorem taxes (i.e., how much less cola is bought) should not outweigh the one to excise taxes. This may be a reasonable assumption for products that are cheap, as an ad-valorem tax is often leading to a lower price increase than taxing sugar content.[27] Therefore, it is unlikely that ad-valorem taxes would generate a larger reduction in the number of (cheap) colas purchased than excise taxes. For other products, this is more ambiguous. However, as long as the price change generated by these two types of taxes is similar for more expensive products, we should not expect major differences in the extensive margin reaction that could invalidate the results of above. Another critical assumption is that the only outside goods in these simulations are diet colas. Excise taxes seems more effective than ad-valorem taxes in making households switch to a diet version. This result can in principle change if we study substitution to other goods. However, following the same reasoning of above, as excise taxes tend to increase the price of cheaper colas by more, they should be also more effective in leading households to substitute towards other product than diet colas.

The simulations above assume that the pass-through of these taxes to prices is equal to 100%. That is, any tax would increase cola prices by the same amount. This assumption is in line with recent evidence from the UK soda tax implemented in 2018, which was perfectly shifted to soda prices (Dubois et al. 2020). However, it cannot be ruled out that firms react more strategically to the implementation of sugar taxes. For instance, Bonnet & Requillart (2013) suggest that excise taxes are usually over-shifted to soft drink prices, while ad-valorem taxes are instead often under-shifted. This could reinforce the effect of above, thus supporting the view that specific taxes on sugar should be preferred on public health grounds. Furthermore, Ally et al. (2014) show that alcohol taxes are typically under-shifted to the price of cheaper alcoholic beverages. Suggesting that the substitution towards cheaper sugary products with ad-valorem taxes may be larger than the one estimated. Interestingly, a possible explanation of the lower tax shifting for cheaper products can be found in the shape of their demand curves. In the context of imperfectly competitive markets, tax pass-through should increase with the convexity of the demand curve (Weyl & Fabinger, 2013; Hindriks & Myles, 2013). The estimated demand curves for the two cheapest colas (D.1 and D.2) display a much lower degree

[27] This is because the ad-valorem tax is levied as a percentage of the product price, which is very low for discount products. For instance, let us take two tax schemes that would rise the average consumer price of sugary colas similarly. Under a 20% ad-valorem tax, the price increase for the discount cola is €0.09, while this is equivalent to €0.33 under an excise tax of €0.20/100gr of sugar.

of convexity compared to more expensive products (see in the appendix). This finding implies that ad-valorem taxes could be shifted more to the price of expensive brands than to the price of cheaper colas.

2.5.4 Model Validation

A possible concern is that the estimates of taste distributions over-fit the subsample in which the model is estimated, thus providing results that are not generalizable to the rest of households. The comparison of the Logit models in Table 2.A.1 suggests that the Logit parameter estimates do not change much with the reduction in sample size. However, this cannot be confirmed for the Mixed Logit, as it is practically unfeasible to estimate this model on the entire sample. One way to validate the results of this model is to compare its in-sample predictive performance against its out-of-sample one. Although cola choice prediction is not the ultimate scope of this study, testing the predictive performance of the model can serve as an indicator of its relevance in explaining household choices. Comparing the predictive performance within and outside the estimation sample can also indicate whether the model is still relevant in explaining the choices of households that were not considered in the estimation procedure. Furthermore, as demand curves represent the outcome of predicting the average purchase probability under different price levels, this exercise can give an idea of how reliable these counterfactual simulations are. The predictive performance of the model is computed in terms of the average percentage error in predicting actual market shares. More formally, the average of the correct predicted shares with the Mixed Logit estimates is given by the following formula:

$$1 - \left[\frac{1}{J} \sum_{j} \left(\frac{|\tilde{s}_j - s_j|}{s_j} \right) \right].$$

Where s_j and $\tilde{s}_j$ are the actual market share and the predicted market share of cola j, respectively.[28] While J is the total number of cola products. The expression in square brackets is equivalent to the percentage prediction error of the model on average. The model is fitted to four distinct samples, where one is the original estimation sample, and the other three are other randomly selected samples of 10.000 households. Fitting the model to its estimation sample, correctly predicts the 79.29% of market shares on average. Interestingly, fitting the model to

[28] Here market shares are considered as the share of shopping occasions in which households choose a given cola. Thus, this is equivalent to $s_j = \frac{\sum_i \sum_t y_{ijt}}{N}$. Where N is the total number of shopping trips in the sample and y_{ijt} is a variable equal to 1 if household i bought cola j during shopping trip t.

the other random samples delivers very similar results, with an average of 78.86% correct share predictions. This result indicates that the model predicts reasonably well actual market shares, even when considering different households and under different market scenarios. Thus, suggesting that the parameter estimates do not over-fit the estimation sample.

2.6 Observed Taste Heterogeneity

The results shown in the previous section suggest that households have very heterogeneous preferences for colas and imply that their behavioral response to sugar taxes can be highly heterogeneous. Accounting for observed taste heterogeneity, however, is fundamental to understand the impact of these policies among those households overconsuming sugary products. Similarly, policymakers may also be interested in understanding how the tax burden and health benefits are distributed across families with different socioeconomic statuses. In order to address these questions, the utility specification outlined in Eq. 2.1 is extended by allowing price sensitivity, taste for sugar and bottle size to vary with observable households' characteristics. Formally, their respective coefficients β_i, δ_i and π_i are replaced in the utility specification 4.1 by:

$$\beta_i = \bar{\beta} + \sum\nolimits_H \beta_H \, (Demo_H)_i + \beta_S (Sugar_{share})_{it} + \zeta_i \sigma_\beta \; ;$$

$$\delta_i = \bar{\delta} + \sum\nolimits_H \delta_H \, (Demo_H)_i + \delta_S (Sugar_{share})_{it} + \eta_i \sigma_\delta \; ;$$

$$\pi_i = \bar{\pi} + \sum\nolimits_H \pi_H \, (Demo_H)_i + \xi_i \sigma_\pi \; .$$

Where the variables $Demo_H$ are the household demographics described in the data section. Which include income, number of household members, a dummy variable indicating whether they have kids, educational level, and age of the head of household. The coefficients β_H, δ_H, and π_H indicate how price sensitivity, taste for sugar and preference for bottle's size change with the observed household demographics. In addition to these household characteristics, price sensitivity β_i and taste for sugar δ_i are allowed to vary with the variable $Sugar_{share}$. Which is the share of sugary dense products in the overall food shopping basket of household i. Their associated coefficients β_S and δ_S indicate how price sensitivity and taste for sugar change with the level of consumption of sugary products.

These two coefficients are crucial in order to evaluate the effectiveness of sugar taxes on public health and internality grounds. If $\delta_S > 0$, heavy sugar consumers tend to prefer sugary to diet colas, and therefore, there is a public health rationale in taxing these products. Conversely, if

$\delta_S \leq 0$, taxing sugary colas may not lead to significant public health benefits, as these products are not particularly appealing to those overconsuming sugar. Sign and magnitude of β_S indicate how effective this policy will be in reducing the demand for sugary colas among the targeted population of heavy sugar consumers. With $\beta_S > 0$, price sensitivity is higher among those overconsuming sugar. Hence, a sugar tax would be quite effective in reducing their demand for sugary colas and so their sugar intake. The larger β_S, the higher the effectiveness of this policy in correcting the household behavior. The variables $\bar{\beta}, \bar{\delta}$ and $\bar{\pi}$ are the intercepts of the taste coefficients and therefore represent the part of preferences that do not change with observable demographics. While the terms ζ_i and η_i are random variables, which are distributed as follows: $\zeta_i \sim TN_{[-\mu_\beta,+\infty]}(0,1)$; $\eta_i \sim N(0,1)$; $\xi_i \sim N(0,1)$. With μ_β representing the mean of β_i, which ensures that $\beta_i \geq 0, \forall i$. These random terms are assumed to be independent to ease the estimation procedure.[29] However, correlation in taste coefficients is allowed through households' demographics.

This model is estimated on a random subsample of households spending at least 80% of their food-at-home budget in this retail chain. This approach allows circumventing the issue of not capturing sugar eating habits of those households purchasing food in different supermarket chains. The total number of households corresponding to this characteristic is 130,514. The Mixed Logit with observed group heterogeneity is estimated on a random subsample of 10,000 households, for a total of 358,225 shopping trips. The results of this estimation are shown in Table 2.8 below.[30] To check for the representativeness of this sample of customers with respect to the overall population, the model of Eq. 2.1 (unobserved heterogeneity) is also estimated on this sample and results are compared to those exposed in Table 2.4. The resulting coefficients are in table 2.A.5 in the appendix, which highlights very similar demand parameters between frequent and all shoppers of this retail chain.

The results displayed in Table 2.8 suggest that price sensitivity, taste for sugar, and preference for bottle's size do indeed vary with observable household characteristics. All coefficients of the interaction variables between taste parameters and household demographics are significant at the 1% level. The sign of these coefficients indicates how demographics are correlated with

[29] It is also assumed that $\beta, \delta, \pi \perp X|Z$. Where X is the set of explanatory variables and Z is the set of households' demographics. This assumption implies that households' preferences are exogenous to product attributes given the households' demographics.

[30] Table 2.A.4 in the appendix shows the parameter estimates of a simple Logit version of this model, which is instead estimated on the entire sample of assiduous costumers. Such a model does not allow for unobserved taste heterogeneity and hence it assumes that $\sigma_\beta = \sigma_\delta = \sigma_\pi = 0$.

taste parameters. As many demographics are included, this has to be interpreted as a partial correlation between the household characteristic and taste coefficients. That is, it is a measure of how taste parameters vary with a given demographic variable while holding all other household characteristics as fixed.

Table 2.8: Parameter estimates of *Dynamic* Mixed Logit with observed heterogeneity

Parameter	Estimate	Demographics	Estimate		
			Price	**Sugary**	**Size**
Price $\bar{\beta}$	**33.268** (0.793)	**log(Income)**	**-3.242** (0.079)	**-1.080** (0.035)	**0.463** (0.057)
Price σ_{β}	**4.253** (0.032)	**Education**	**-0.193** (0.077)	**-0.330** (0.031)	**-0.786** (0.052)
Sugary $\bar{\delta}$	**10.222** (0.351)	**Sugar share**	**-8.222** (0.312)	**6.273** (0.152)	
Sugary σ_{δ}	**2.347** (0.014)	**Household's size**	**0.254** (0.012)	**0.119** (0.006)	**0.263** (0.009)
Size $\bar{\pi}$	**-4.911** (0.572)	**Kids**	**-0.343** (0.038)	**0.006** (0.017)	**-0.275** (0.027)
Size σ_{π}	**2.835** (0.014)	**Age>70**	**-0.517** (0.080)	**0.384** (0.034)	**-1.534** (0.054)
Rare $\bar{\theta}$	**-0.972** (0.011)	**70>Age>58**	**0.659** (0.066)	**0.173** (0.026)	**-1.327** (0.044)
Folder $\bar{\omega}$	**-0.036** (0.019)	**58>Age>46**	**1.091** (0.055)	**0.475** (0.021)	**-0.935** (0.035)
Rebate $\bar{\varphi}$	**0.173** (0.018)	**46>Age>34**	**0.586** (0.053)	**0.101** (0.020)	**-0.529** (0.034)
State dependence $\bar{\gamma}$	**1.930** (0.004)				
Brand fixed effects			Yes		
N° of Observations			358,225		
McFadden's R^2			0.57		

Notes: *standard errors are in parenthesis. The first and second columns show intercepts and standard deviations of random taste parameters. Together with the estimates of the fixed taste parameters. From the third to the last column, the table displays the estimates of the interactions between price, sugary and size coefficients with observed demographics. Age coefficients are computed with respect to families with a head of household below 34 years old. All parameters are significant at the 1% level except for the interaction between kids and the taste for sugar.*

The parameters displayed in Table 2.8 indicate that taste coefficients do vary with household demographics. At the same time, their estimated standard deviations also suggest that a substantial part of such heterogeneity is still not explained by differences in observed household characteristics. This finding suggests that a Mixed Logit model should be preferred to a simple Logit, as household preferences also vary randomly across the population. A possible concern for the Mixed Logit model with observed household demographics is that the reduction in sample size can undermine the correct identification of observed taste heterogeneity across households. However, this does not seem to be the case for these estimates. Most of the interaction coefficients are significant at the 1% level. Furthermore, the simple Logit version of the model, which was estimated on the entire sample of assiduous costumers, shows very similar results in terms of observed taste heterogeneity across households.

Price sensitivity is decreasing with income and educational level, and it is lower among households having kids, while it increases with family size. These results are somehow expected, as poorer, larger, and low-educated households tend to be more resource-constrained and thus more sensitive to prices. For households having similar characteristics, old families (with the head of household above 70 years old) tend to be less price-sensitive, followed by the youngest (with the head of household below 34 years old). Interestingly, the estimates suggest that price sensitivity decreases with the household share of sugary-dense products in the overall shopping basket. This means that households consuming more sugar are less sensitive to changes in cola price compared to healthier households.

Taste for sugary colas also tends to vary with observable household characteristics. Poorer, low-educated, and larger households prefer sugary to diet colas. The relationship between age and preference for sugary colas is instead ambiguous. Oldest families and those with the head of household between 46 and 58 years old are more likely to prefer a sugary cola compared to those in other age groups. Not surprisingly, households that usually consume more sugar have stronger preferences for sugary colas. The higher the total consumption of sugar (net of that one consumed in colas), the higher the probability of consuming a sugary cola. This result suggests that households perceive sugar in colas as a sort of complement to the overall sugar consumption. A household in the lower quartile of the sugar consumption distribution (i.e., the share of sugary-dense products) has a probability of 41.25% of choosing a sugary cola. In contrast, a household in the top quartile of the sugar consumption distribution has a 49.65% probability of choosing a sugary cola. Table 2.8 also indicates that higher-income households prefer larger bottle sizes, while highly educated families tend to prefer smaller bottles. Larger families are more likely to go for large bottles, although this effect is smaller if the household has kids. Contrary to both price and sugary parameters, the relationship between age and size coefficients is somehow linear. The older the family, the higher the preference for smaller bottle sizes.

2.6.1 The Impact of Taxation across Household Types

The estimates of the *dynamic* Mixed Logit model with observed group heterogeneity allows evaluating the impact of sugar taxes on cola demand across households with different demographics and sugar eating habits. The results of this model highlight two important empirical facts. First, heavy sugar consumers tend to prefer sugary to diet colas. This indicates that there is a public health rationale in taxing these products, as their consumption is more prevalent among those households that have higher health risks. Second, these households are not more sensitive to price changes than others. Thus, suggesting that sugar taxes would not be more effective in reducing cola demand among the targeted population of heavy sugar consumers. However, the model estimates also suggest that poorer households are likely to benefit the most in terms of public health. This because they have both a greater preference for sugary colas and higher price sensitivity compared to richer households. Therefore, a sugar tax would be quite effective in inducing these households to switch from sugary to diet colas. The fact that poor households are very price-sensitive also helps to reduce their tax burden in monetary terms. In a broader economic welfare framework, these findings can suggest that a sugar tax on colas may be both efficient and progressive (see Griffith et al., 2018; Allcott et al., 2019). A necessary condition, however, is that sugar internalities exist and decrease with the level of income.

Table 2.9 below shows the potential impact of an excise tax of €0.30 per 100 grams of sugar across different household groups. Lower-income households, who generally prefer sugary colas, are more likely to switch to a diet cola since they are more price-sensitive than higher-income households. Conversely, high-income households, which generally tend to dislike sugary colas, reduce their purchase probability by less than poorer households. In particular, with the implementation of this tax, the probability of choosing a sugary cola decreases by -24.24% for low-income households in the long-run, against the -19.24% for high-income households. Moreover, this tax is less effective in reducing sugary cola demand among heavy sugar consumers than among healthier households. This is because frequent consumers of sugar are less sensitive to cola prices, even though they have a higher demand for sugary products.

The estimates in Table 2.9 shows that the implementation of the tax reduces the demand for sugary colas by -18.57% among heavy sugar consumers in the long-run, while this reduction amounts to -24.46% for low sugar consumers. The impact of this tax, however, is quite similar across households with different educational levels. In terms of tax incidence, heavy sugar consumers and poor households are the groups bearing this tax the most.[31] This effect, however, is smaller for poor households, as they are more sensitive to prices and so they switch more quickly to untaxed diet products. As for the tax simulations displayed in Table 2.7, this exercise suggests that a sugar tax would be slightly more effective in the long-run due to state dependence in cola choice. Since it is assumed for simplicity that state dependence is fixed across households, this dynamic effect does not vary much across different groups. In Subsection 2.6.3, this hypothesis is relaxed to investigate some possible heterogeneity in state dependence.

Table 2.9: Impact of €0.30 excise tax per 100gr sugar across household types

Household types	**Demand for sugary colas (average purchase probability)**				
		Short run		**Long run**	
	Before	**After**	**Δ%**	**After**	**Δ%**
Low Income (Q1)	48.01%	37.21%	**-22.50%**	36.37%	**-24.24%**
High Income (Q4)	42.42%	34.95%	**-17.61%**	34.13%	**-19.54%**
Low Sugar Consumption (Q1)	41.53%	32.17%	**-22.54%**	31.37%	**-24.46%**
High Sugar Consumption (Q4)	49.65%	41.30%	**-16.82%**	40.43%	**-18.57%**
Low Education (Q1)	45.60%	36.60%	**-19.74%**	35.76%	**-21.58%**
High Education (Q4)	43.74%	35.08%	**-19.80%**	34.26%	**-21.67%**

Notes: *in parenthesis, there are the quartiles of the distribution for each household type. Q1 stands for the first quartile, while Q4 is the last quartile. The column "before" and "after" indicates the market share of sugary colas before and after the tax simulation.*

In Subsection 2.5.3, it was shown that excise taxes should be preferred to ad-valorem taxes because they can prevent households from substituting to cheaper sugary colas. Accounting for observed taste heterogeneity, however, allows investigating the demographic composition of these households. In particular, it is interesting to understand whether these households are part of the targeted population of heavy sugar consumers. This finding would undermine the effectiveness of ad-valorem taxes on both public health and internalities motives. It is also crucial to check whether these are low-income households, which would then suggest that ad-valorem taxes tend to be more regressive than excise taxes.

[31] This result is only partial as the intensive margin is not considered here. Poor households and heavy sugar consumers can be still more likely to purchase the taxed cola but reduce their number of colas bought more than other households. This, in turn, would decrease their total tax burden.

These questions are addressed by simulating the impact of a 30% ad-valorem tax on the purchase probability of two distinct categories of consumers. That are: (i) low-income households with a high sugar consumption (i.e., those in both the first and last quartile of the income and sugar consumption distribution, respectively); (ii) high-income households with a low sugar consumption (i.e., those in both the last and first quartile of the income and sugar consumption distribution, respectively). Table 2.10 shows the results of this simulation by presenting the change in the purchase probability of both the cheapest and the other more expensive sugary colas for these two categories of households. The simulation of a 30% ad-valorem tax highlights important differences in the behavioral response across different household types. Low-income households with high sugar consumption are much more likely to substitute for a cheaper sugary cola in response to the ad-valorem tax. In particular, their demand for the cheapest sugary cola increases by 15.43%, while this increase amounts to only 4.11% for high-income households with low sugar consumption.

Table 2.10: Impact of 30% ad-valorem tax across household types

Household types	Cheapest sugary cola			Other sugary colas			All sugary colas		
	Before	**After**	**Δ%**	**Before**	**After**	**Δ%**	**Before**	**After**	**Δ%**
Low Income & High Sugar Consumption	4,66%	5,38%	**+15,4%**	49,45%	40,43%	**-18,2%**	54,10%	45,81%	**-15,3%**
High Income & Low Sugar Consumption	2,76%	2,88%	**+4,1%**	35,81%	29,01%	**-19%**	38,57%	31,89%	**-17,3%**

Notes: *The columns "before" and "after" indicate the market share of each type of cola before and after the tax simulation. The cheapest sugary cola is brand D.1, while the other sugary colas are all the remaining sugary brands. The column "All sugary colas" refers to the overall market share for sugary colas, which is the sum of the cheapest and the other sugary brands. The table just report the long-run impact of the tax.*

This substitution effect makes the ad-valorem tax less effective in reducing demand for sugary colas, especially among the targeted population of heavy sugar consumers and low-income households. On public health grounds, this finding reinforces the hypothesis that excise taxes should be preferred to ad-valorem taxes in order to reduce sugar consumption among the targeted population. Furthermore, since the consumption of cheap sugary colas is more prevalent among the poor, this also suggests that ad-valorem taxes could be regressive. With an excise tax, the demand for cheap sugary colas would decline rapidly since the incentive to substitute for a diet cola is higher than with an ad-valorem tax (see Figure 2.2.b). Although the demand for cheaper sugary colas is more prevalent among the poor, this larger substitution towards diet colas reduces their tax burden in monetary terms and makes this tax also more effective in reducing their sugar intake.

2.6.2 Households' versus Individuals' Preferences

One possible concern when studying households' preferences is the aggregation of taste across different individuals. This paper assumes that taste for bottled colas are homogeneous across different individuals within the household. This assumption may be plausible given the nature of the product considered. While colas in cans are made for individual consumption, bottled colas (above 1 liter) can be consumed by more individuals. This seems to be confirmed also when observing the data in this sample. Households chose more than one brand of cola in a single shopping trip only 5% of the times, thus suggesting that people chose only one type of cola to be consumed in the family.

Table 2.11: Households versus Individuals

Parameter	***Individuals***	***Households***
Price $\bar{\beta}$	**2.531**	**2.525**
	(0.185)	(0.040)
Price σ_β	**4.585**	**4.417**
	(0.156)	(0.033)
Sugary $\bar{\delta}$	**-0.004**	**0.519**
	(0.089)	(0.023)
Sugary σ_δ	**2.632**	**2.345**
	(0.073)	(0.015)
Size $\bar{\pi}$	**-0.617**	**-0.407**
	(0.080)	(0.017)
Size σ_π	**2.862**	**2.821**
	(0.069)	(0.014)
Rare $\bar{\theta}$	**-0.938**	**-0.957**
	(0.052)	(0.011)
Folder $\bar{\omega}$	**0.020**	**-0.029**
	(0.041)	(0.009)
Rebate $\bar{\varphi}$	**0.135**	**0.176**
	(0.038)	(0.008)
State Dependence $\bar{\gamma}$	**1.932**	**1.932**
	(0.020)	(0.004)
Brand Fixed Effects	Yes	Yes
Observations	16,971	341,254
McFadden's R^2	0.52	0.54

Notes: *standard errors are in parenthesis.*

However, it cannot be ruled out that the estimation approach undertaken in this paper conceals some additional intra-household heterogeneity in taste. Having data about households' composition allows exploring how taste can vary between households and individuals. For households composed of a single member, in fact, such issue of taste aggregation does not arise. Estimating the cola choice model on these individuals separately allows detecting their preferences quite precisely. Table 2.11 above reports the results of this estimation. The model is estimated allowing for state dependence and unobserved taste heterogeneity, and it is estimated on the sample of frequent costumers. Interestingly, the results displayed in Table 2.11 show very similar taste coefficients for households and individuals. Patterns of heterogeneity and state dependence do not vary much across these two groups. The only difference is in the average taste for sugary colas, which tend to be higher among households.[32] This is also consistent with table 2.8, where the taste for sugar increases with the household size. This similarity in demand estimates suggests that households tend to behave just as individuals. Hence, although some intra-household heterogeneity can still exist, this should not invalidate the demand estimates substantially.

2.6.3 Heterogeneity in State Dependence

In all models above, taste heterogeneity across households was allowed only for the three main product characteristics, such as price, sugar, and size. The choice of these three variables is justified by the fact that imposing random coefficients on these product feature can still deliver realistic substitution patterns while limiting the computational burden. To reduce the computational burden, however, it is also possible to decrease the sample size further. This may not be optimal for a correct estimation of taste heterogeneity, as it limits the number of observations for the different households' type. Nevertheless, this allows exploring heterogeneity in other utility parameters, notably the one for state dependence.

So far, it has been assumed that state dependence was constant in the population. This means that all households have the same propensity in doing persistent choices driven by their recent shopping history. This assumption may be incorrect as this behavior can be concentrated just on one particular type of household. To investigate heterogeneity in state dependence, both models with observed and unobserved heterogeneity in taste and state dependence are re-estimated by allowing heterogeneity in the state dependence coefficient. In particular, the state dependence parameter is allowed to be normally distributed (for simplicity it is assumed no correlation across taste parameters). The sample is reduced to 1000 households to ease the estimation procedure.[33] The results of these estimations are displayed in Table 2.12 below.

Table 2.12 indicates that the magnitude of state dependence can vary across households. The mean coefficient $\bar{\gamma}$ is like the one displayed in Table 2.8, where it was assumed a fixed state dependence parameter. However, its standard deviation is positive and significant, meaning that

[32] Standard errors are smaller in the households' model as the sample is quite larger.

[33] However, similar results are obtained with other random subsamples of larger size.

households have a different propensity in doing state-dependent choices. Most households exhibit state dependence in cola choice as the 95% of households have a positive state dependence parameter. This means that their probability of choosing a given cola increases if the same cola was purchased in the previous shopping trip. The rest of households have instead slightly negative state dependence parameters, which suggest either no state-dependent choices or variety-seeking behavior (i.e., the probability of choosing a cola decreases if this cola was purchased in the previous shopping trip). The magnitude of state dependence tends to vary also across observed demographics. In particular, it is greater for young families (head of household below 45), larger and more-educated households, and heavy sugar consumers. It is smaller for high-income households and families without kids.[34]

Table 2.12: Heterogeneity in State Dependence

Unobserved Heterogeneity		**Observed and Unobserved Heterogeneity**	
Parameters	**Estimate**	**Parameters**	**Estimate**
State dependence $\bar{\gamma}$	**1.978** (0.022)	**State dependence $\bar{\gamma}$**	**5.501** (1.020)
State dependence σ_γ	**1.189** (0.029)	**State dependence σ_γ**	**1.405** (0.022)
		log(Income)	**-0.337** (0.102)
		Education	**1.023** (0.081)
		Sugar share	**1.121** (0.081)
		Household's size	**0.056** (0.016)
		Kids	**-0.714** (0.046)
		Age>70	**-0.287** (0.011)
		70>Age>58	**-0.587** (0.011)
		58>Age>46	**-0.464** (0.090)
		46>Age>34	**0.132** (0.090)
Brand fixed effects		Yes	
N° of Observations		30,534	
McFadden's R^2		0.56	

Notes: *standard errors are in parenthesis.*

[34] While these descriptive figures can be useful to analyze how state dependence can vary across different socio-economic groups, they should be taken with caution as the model is estimated on a smaller sample size. Yet, re-estimating the model on other random subsample of equal size delivers similar results.

The previous sections show that state dependence implies a dynamic demand, where the probability of switching from a sugary to a diet cola due to a tax hike increases in the long-run. Therefore, the heterogeneity in state dependence across households can have important implication for tax policy. The results in Table 2.12 indicates that this effect can be greater for low-income households and heavy sugar consumers. This suggests that a sugar tax will be initially less effective for these households. Nevertheless, their probability of switching to a sugary cola will increase with time and their tax burden will be eventually reduced. This dynamic can be important when evaluating the short-run impact of an actual tax reform. Indeed, a short-run reduced-form evaluation of a sugar tax could understate the long-run impact of this policy among the targeted population of low-income households and heavy sugar consumers.

2.7 Robustness Checks

The empirical evidence presented so far suggests that households have heterogeneous tastes and positive state dependence in product choice. In the previous sections, it was also discussed the importance of disentangling these two behaviors and their implications in terms of sugar tax design. This section of the paper presents the results of various robustness checks that were performed to confirm these findings and rule out other possible explanations for the observed inertia in cola choice. The two main hypotheses to be rejected are that the estimated state dependence is the result of either a misspecification of taste heterogeneity distribution or autocorrelation in the error terms.

If the distribution of taste heterogeneity is not correctly specified, the coefficient on the lagged choice variable may still be positive because it captures some persistent propensity of the household to make a given choice. An example of this effect can be seen in table 2.4. When accounting for taste heterogeneity through random coefficients, the state dependence parameter dropped considerably compared to the logit specification (see column 2 against column 4). Hence, it should be checked whether accounting for a more flexible distribution of taste heterogeneity across the population would still deliver a positive state dependence coefficient.

Following Dubé et al. (2010), this is tested by randomly reshuffling the order of the households' shopping trips and then checking whether the lagged choice coefficient is still positive and significant. The rationale behind this test is the following. If the estimated state dependence is due to some misspecification of taste heterogeneity, then changing the order of households' shopping trips should still deliver a positive state dependence coefficient. In contrast, if taste heterogeneity is correctly specified, only the recent shopping experience should influence the current choice. Hence, changing the order of the shopping trips should deliver a coefficient on the lagged choice variable that is close or equal to zero.

The result of this analysis is displayed in Table 2.13 below. The row *actual* shows the coefficient on the lagged choice variable when the reshuffled lagged choice is equal to the actual lagged choice. The row *reshuffled* shows the lagged choice coefficient when the reshuffled lagged choice is not equal to the actual lagged choice.[35] The test is performed under the two different specifications of household taste heterogeneity. The first column shows the results of this test under unobserved taste heterogeneity. The second column shows the results of this test under both observed and unobserved taste heterogeneity.

Table 2.13: State Dependence with Random Reshuffling

	Random Reshuffling	
	Unobserved Heterogeneity	***Unobserved Heterogeneity with HH Demographics***
Actual	**2.121** (0.007)	**2.167** (0.006)
reshuffled	**0.339** (0.008)	**0.252** (0.007)

Notes: *standard errors are in parenthesis. All coefficients are significant at the 1% level.*

The reported coefficients suggest that the hypothesis that state dependence is due to a misspecification of taste heterogeneity should be rejected. The coefficient of the actual lagged choice is positive and significant, while that of the reshuffled lagged choice is close to zero. The fact that the latter is positive and significant, however, does suggest the existence of some misspecification in the distribution of taste heterogeneity. However, this seems too small to explain the observed state dependence in cola choice. Accounting for a more flexible distribution of taste through household demographics reduces this misspecification sensibly. The coefficient on the reshuffled choice goes in fact closer to zero.

To check whether state dependence is robust to different specifications of taste heterogeneity, the basic model specification in Eq. 2.1 is re-estimated by allowing the state dependence parameter γ to be randomly distributed and by specifying a more flexible distribution of taste across households. Given their high computational burden, these models are estimated on a

[35] By randomly reshuffling the order of the shopping trips, around 50% of reshuffled lagged choices are still equal to the actual past choice. Although here it is presented the result of one random reshuffling only, this exercise was performed many times with different reshuffled sequences and always delivered similar results.

random subsample of 1000 households. One model is estimated with a normal distribution of the brand effects as well as for price, sugary and size parameters. The distribution of state dependence parameter for this model is displayed by the dashed line in Figure 2.3. Furthermore, another model is estimated where price, sugary, size, and state dependence parameters are allowed to be distributed as a mixture of three normal distributions. The distribution of the state dependence parameters for this model is shown by the red line in Figure 2.3. As shown in Figure 2.3, going from the most basic specification of heterogeneity (black line) to the more complex ones (dashed and red lines), slightly reduce the magnitude of the state dependence parameter.[36] However, this remains positive and significant for most households.

Figure 2.3: State Dependence Coefficients for Different Heterogeneity Specifications

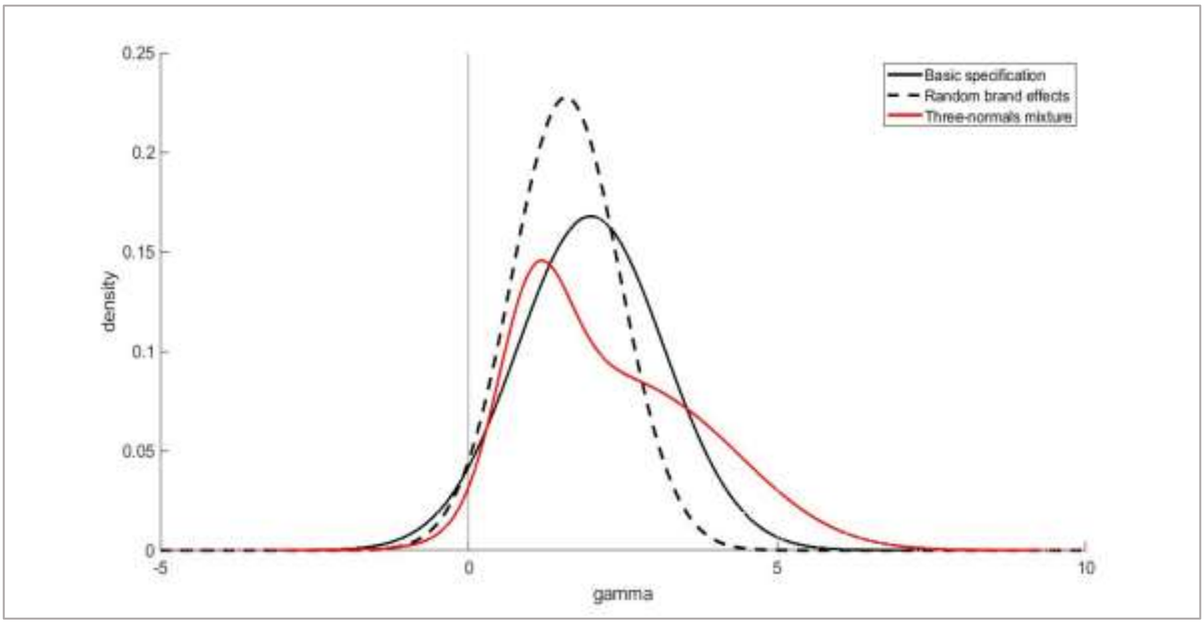

A model with true state dependence and model with autocorrelated error terms have somewhat different policy implications. If unobservable factors drive choice dynamics, governments cannot exploit state dependence to enhance the impact of sugar taxes in the long-run. Three different tests are performed in order to rule out this hypothesis. Their results are presented in Table 2.14, and each of them is discussed in turn below. The first test follows Chamberlain (1985). It is estimated a model where the lagged choice variable is substituted by lagged prices (i.e., the price vector faced by the household in the previous shopping trip). If choice inertia is due to autocorrelated error terms, then past prices should not influence the current choice probabilities. Conversely, in case of state dependence, lagged prices should negatively affect

[36] The black line shows the distribution of the state dependence parameter where price, sugary, and size variables are distributed as in the basic Mixed Logit specification. The parameters of this distribution are also those displayed in Table 2.12.

current choice probabilities as they can influence the household's state variable. This means that the higher is the lagged price of cola, the lower should be its current purchase probability.

Table 2.14: Robustness Checks for Autocorrelation

Test for Autocorrelation					
Lagged Price		***State Dependence with Price Discount***		***Cola not Available in Last Choice Set***	
$price_{t-1}$	**0.694*****	SD	**1.927*****	NA_{t-1}	**-0.413*****
	(0.041)		(0.007)		(0.032)
		$SD * D_{t-1}$	**-0.020**		
			(0.011)		

Notes: *standard errors are in parenthesis. The *** indicates statistical significance at the 1% level. All coefficients are obtained by taking as baseline model the dynamic Mixed Logit without household demographics. Similar results are obtained by including household demographics.*

The result of this test is shown in the first column of Table 2.14 above, where the lagged price coefficient is shown in the first column next to the row labelled $price_{t-1}$. The hypothesis of autocorrelated error terms is rejected as the lagged price coefficient is positive and significant (remember that price enters the utility specification as a negative term). Higher past prices reduce the current purchase probabilities, hence proving evidence of positive state dependence. The main issue with this test, however, is that lagged prices might proxy for households' price expectations (Chamberlain, 1985; Dubé et al., 2010). This means that lagged prices can potentially influence current household choices even in the absence of state dependence.

Following the suggestion of Dubé et al. (2010), it is performed a further test for autocorrelation in the error terms by checking at the impact of temporary price cuts on the measured state dependence. They show that in a model with autocorrelated error terms the estimated coefficient for state dependence should be smaller when a price discount initiated the past state. The idea behind this test is that the past error term is larger when the past choice is made at a regular price rather than at a price discount. Hence, if the lagged choice variable is capturing some autocorrelation in the error terms, then the state dependence coefficient must be larger when the past choice was made at a regular price. This is because the lagged choice proxies for both a large past and current random utility draw. This test is performed by adding to the utility specification an interaction term between the lagged choice variable and a dummy variable indicating whether the product was under discount in the previous shopping trip. A product was considered under discount if the price in the previous shopping trip was at least 5% lower than in the current shopping trip.

The result of this test is displayed in the second column of Table 2.14. The variable SD indicates the state dependence coefficient on the lagged choice variable. While the interaction term $SD * D_{t-1}$ measures how state dependence varies if the product was under discount in the previous shopping trip.[37] The hypothesis of autocorrelated error terms is also rejected by this test. The coefficient of the interaction variable does not capture any lower state dependence when a previous price discount initiated the past state.

The nature of the dataset allows proposing a further test for autocorrelated error terms. The lagged choice variable is substituted with a dummy variable indicating whether a product was not available in the previous shopping occasion. In a model with autocorrelated error terms, a product that was not available in the previous shopping trip should not have any impact on the current choice probability. In a model with true state dependence, however, this dummy variable can affect current choice probabilities as a change in the household's previous choice set can influence its state variable. In particular, the coefficient of this dummy variable should be negative and significant since the absence of a product in the previous choice set prevents this product from being purchased in the past, and so reduces its current choice probability. The third column of Table 2.14 shows the result of this test. The variable NA_{t-1} represents the coefficient of the dummy variable that indicates whether a product was not available in the previous shopping occasion. This coefficient is negative and significant. Hence, suggesting the existence of a true state dependence driving the households' choice dynamics.

The model is also re-estimated by including a set of time-varying brand intercepts to control for some possible unobservable factors that might be correlated with prices. This check controls notably for advertising campaigns, which are done nationally and are unobserved in the data. The results of this estimation can be found in Table 2.A.6 in the appendix. As shown in this table, taste coefficients are very similar to the one displayed in Table 2.4 (column 4). Thus, suggesting that time-varying unobservable factors, such as advertising, do not affect much the results.

2.8 Conclusions

Sugar taxes are often considered as a possible tool to tackle excessive sugar consumption. As the issue of excessive sugar intake does not affect all consumers in the same way, accounting for heterogeneity in preferences is crucial to assess the impact of these policies among those overconsuming sugary products. This work estimates a dynamic multinomial Logit model of

---

[37] D_{t-1} is a dummy variable indicating whether the product was under discount in the previous shopping trip.

cola demand on a novel set of supermarket scanner data to study preference heterogeneity across households. A particular focus is given to disentangling persistent taste heterogeneity from other sources of choice inertia, such as state dependence in product choice. The model estimates allow evaluating the effectiveness of taxation in reducing demand for sugary colas across different household types.

The results of this analysis show that households have very heterogeneous preferences for sugary colas and have different price sensitivity. Heavy sugar consumers tend to prefer sugary to diet colas but are less sensitive to cola prices. This finding suggests that although taxing sugary colas can be justified on public health grounds, this policy would not have a greater impact on the targeted population. Poor households, however, are likely to experience larger health gains since they would be more reactive to sugar taxes. Their higher tax responsiveness also helps them reduce their tax burden in monetary terms. The estimated demand model allows performing tax policy simulations that account for very flexible substitution patterns. These simulations show that specific taxes targeting sugar content should be preferred to ad-valorem taxes on sugary colas on both corrective and equity grounds. The main reason for this finding is that ad-valorem taxes would entail a larger substitution towards cheaper sugary brands and that this effect is mostly concentrated among low-income households and heavy sugar consumers.

Lastly, this work finds robust evidence of state dependence in cola choice. This implies that households need some time to fully adjust demand to a price change. Ignoring this behavior can lead to an overestimation of the degree of preference heterogeneity in the population. Because of state dependence, the impact of sugar taxes on household demand would tend to increase over time. The tax policy simulations show that sugar taxes would mostly affect demand in the short-run, but their impact would also have a moderate increase with time.

References

Allais, O., Bertail, P., & Nichèle, V. (2010). The effects of a fat tax on French households' purchases: a nutritional approach. *American Journal of Agricultural Economics*, *92*(1), 228-245.

Allcott, H., Lockwood, B. B., & Taubinsky, D. (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *The Quarterly Journal of Economics*, *134*(3), 1557–1626.

Allenby, G. & Lenk, P. (1994). Modeling Household Purchase Behavior with Logistic Normal Regression. *Journal of the American Statistical Association*, *89*(428), 1-14.

Ally, A. K., Meng, Y., Chakraborty, R., Dobson, P. W., Seaton, J. S., Holmes, J., ... & Meier, P. S. (2014). Alcohol tax pass-through across the product and price range: do retailers treat cheap alcohol differently?. *Addiction*, *109*(12), 1994-2002.

Alvarado, M., Unwin, N., Sharp, S. J., Hambleton, I., Murphy, M. M., Samuels, T. A., Suhrcke, M. & Adams, J. (2019). Assessing the impact of the Barbados sugar-sweetened beverage tax on beverage sales: an observational study. *International Journal of Behavioral Nutrition and Physical Activity*, *16*(1), 13.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242-262.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, *63*(4), 841-890.

Bonnet, C., & Réquillart, V. (2013). Tax incidence with strategic firms in the soft drink market. *Journal of Public Economics*, *106*, 77-88.

Bonnet, C., & Réquillart, V. (2018). *The effects of taxation on the individual consumption of sugar-sweetened beverages* (No. 16-638). Toulouse School of Economics (TSE).

Chamberlain, G. (1985). Heterogeneity, omitted variable bias, and duration dependence. In J. Heckman & B. Singer (Eds.), *Longitudinal Analysis of Labor Market Data* (Econometric Society Monographs, 3-38). Cambridge University Press.

Chouinard, H. H., Davis, D. E., LaFrance, J. T., & Perloff, J. M. (2007). Fat taxes: big money for small change. *In Forum for Health Economics & Policy*, *10*(2).

Coibion, O., Gorodnichenko, Y., & Hong, G. H. (2015). The cyclicality of sales, regular and effective prices: Business cycle and policy implications. *American Economic Review*, *105*(3), 993-1029.

Deaton, A., & Muellbauer, J. (1980). An almost ideal demand system. *American economic review*, *70*(3), 312-326.

Dobbs, R., Sawers, C., Thompson, F., Manyika, J., Woetzel, J. R., Child, P., ... & Spatharou, A. (2014). *Overcoming obesity: an initial economic analysis.* McKinsey global institute.

Dubé, J. P., Hitsch, G. J., & Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, *41*(3), 417-445.

Dubois, P., Griffith, R., & O'Connell, M. (2020). *How well targeted are soda taxes?* (No. 17-868). Toulouse School of Economics (TSE).

Einav, L., Leibtag, E., & Nevo, A. (2010). Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?. *QME, 8*(2), 207-239.

Erdem, T., Keane, M. P., & Sun, B. (2008). A dynamic model of brand choice when price and advertising signal product quality. *Marketing Science*, *27*(6), 1111-1125.

Goeree, M. S. (2008). Limited information and advertising in the US personal computer industry. *Econometrica*, *76*(5), 1017-1074.

Griffith, R., O'Connell, M. & Smith, K. (2018). Corrective taxation and internalities of food consumption. *CESifo Economic Studies*, *64,* 1–14.

Gruber, J., & Kőszegi, B. (2004). Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes. *Journal of Public Economics*, *88*(9-10), 1959-1987.

Harding, M., & Lovenheim, M. (2017). The effect of prices on nutrition: comparing the impact of product-and nutrient-specific taxes. *Journal of health economics, 53*, 53-71.

Heckman, J. (1978). Simple statistical models for discrete panel data developed and applied to tests of the hypothesis of true state dependence against the hypothesis of spurious state dependence. *Annales de l'INSEE 30/31*, 227– 69.

Hendel, I., & Nevo, A. (2006). Sales and consumer inventory. *The RAND Journal of Economics*, *37*(3), 543-561.

Hindriks, J., & Myles, G. D. (2013). *Intermediate public economics*. MIT Press.

Hindriks, J., & Serse, V. (2019). Heterogeneity in the tax pass-through to spirit retail prices: Evidence from Belgium. *Journal of Public Economics, 176,* 142-160.

Jensen, J. D., & Smed, S. (2007). Cost-effective design of economic instruments in nutrition policy. *International Journal of Behavioral Nutrition and Physical Activity*, *4*(1), 10.

Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics*, *15*(3), 310-327.

Levin, L., Lewis, M. S., & Wolak, F. A. (2017). High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, *9*(3), 314-47.

Malik, V. S., Pan, A., Willett, W. C., & Hu, F. B. (2013). Sugar-sweetened beverages and weight gain in children and adults: a systematic review and meta-analysis. *The American journal of clinical nutrition*, *98*(4), 1084-1102.

McFadden, D. (1974), Conditional logit analysis of qualitative choice behavior. P. Zarembka (Ed.), *Frontiers in Econometrics*, Academic Press, New York, 105–142.

Nevo, A. (2000). A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics & Management Strategy, 9*(4), 513-548.

Nordström, J., & Thunström, L. (2009). The impact of tax reforms designed to encourage healthier grain consumption. *Journal of health economics, 28*(3), 622-634.

O'Donoghue, T., & Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, *90*(10-11), 1825-1849.

Seetharaman, P. B., Ainslie, A., & Chintagunta, P. K. (1999). Investigating household state dependence effects across categories. *Journal of Marketing Research*, *36*(4), 488-500.

Sharma, A., Hauck, K., Hollingsworth, B., & Siciliani, L. (2014). The effects of taxing sugar-sweetened beverages across different income groups. *Health economics, 23*(9), 1159-1184.

Te Morenga, L., Mallard, S., & Mann, J. (2013). Dietary sugars and body weight: systematic review and meta-analyses of randomised controlled trials and cohort studies. *Bmj*, *346*, e7492.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Wang, E. Y. (2015). The impact of soda taxes on consumer welfare: implications of storability and taste heterogeneity. *The RAND Journal of Economics*, *46*(2), 409-441.

Weyl, E. G., & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528-583.

World Health Organization. (2017). *Obesity and Overweight Fact Sheet 2018*. Available at < http://www.who.int/en/news-room/fact-sheets/detail/obesity-and-overweight> (Accessed on 6 May 2018).

2.A Appendices

2.1.A: List of food categories used to compute the stock of past sugar consumption

- Biscuits
- Bonbons
- Cakes
- Candies
- Candy bars
- Canned fruits
- Cereals
- Chewing gums
- Choco pasta
- Chocolate bars
- Chocolate beverages
- Chocolate milk
- Chocolate snacks
- Donuts
- Energy drinks
- Fruit juices
- Honey
- Ice creams
- Jams
- Muffins
- Other desserts
- Pastry
- Pudding
- Soya milks
- Sugar sweetened beverages (Cola excluded)
- Syrups
- Table sugar
- Wafers
- Waffles
- Yogurt

2.2.A: Figures and Tables

Figure 2.A.1: Distribution of sugar share

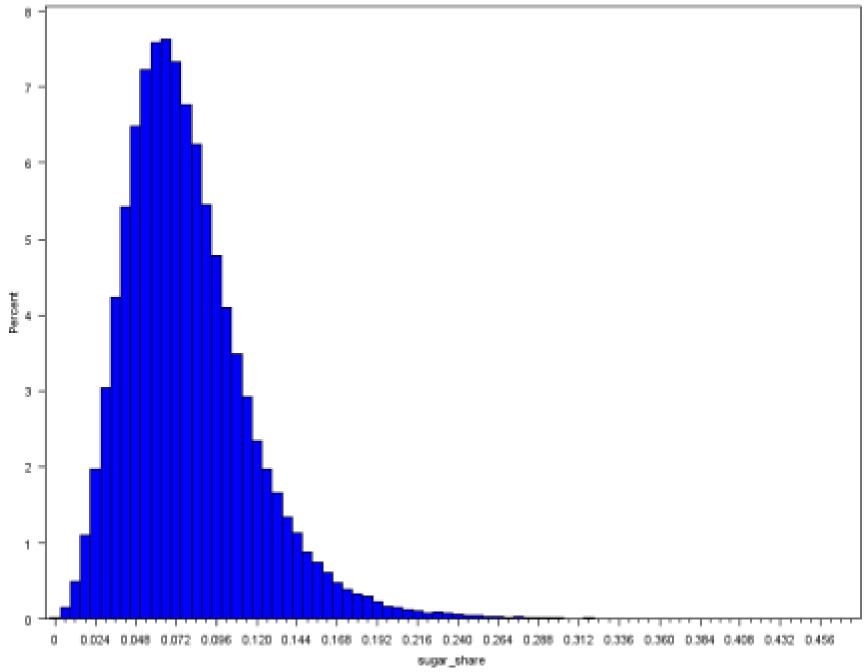

Figure 2.A.2: Spatial price dispersion

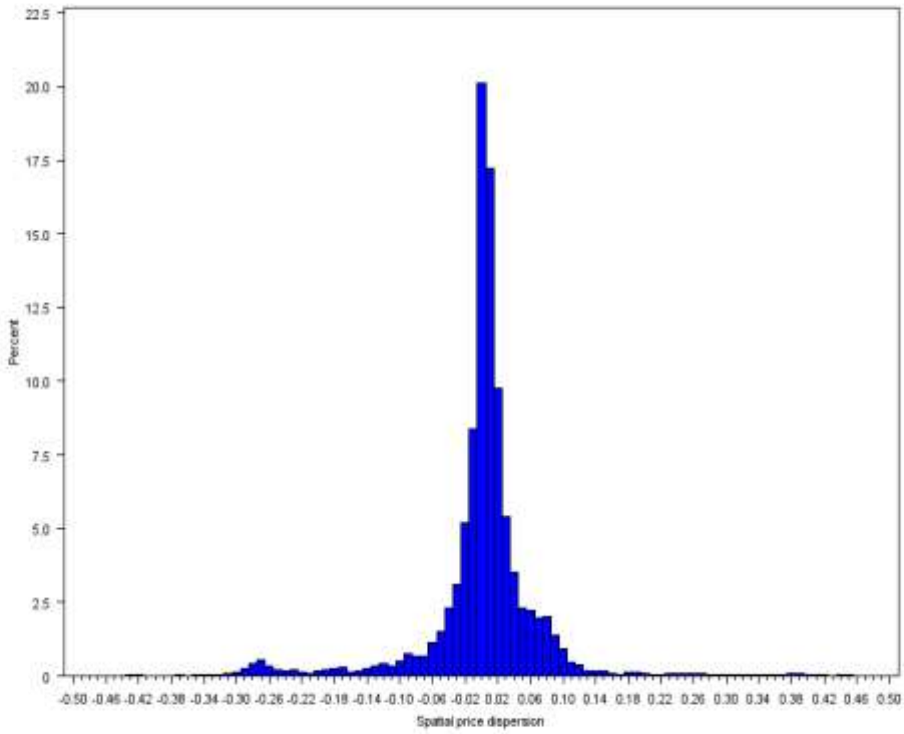

Notes: *price dispersion over stores is computed as the difference between each product price and its daily average across all stores.*

Figure 2.A.3: Temporal price dispersion

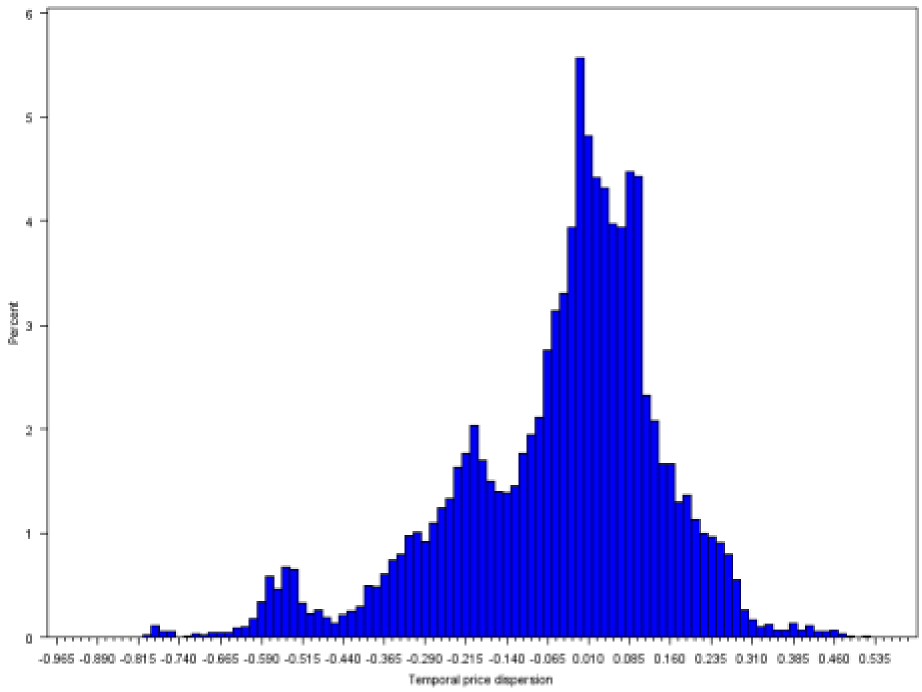

Notes: *price dispersion over time is computed as the difference between the product price and its average in time.*

Table 2.A.1: Logit estimates (full sample)

	Random Subsample (10.000 HHs)		Full Sample (434,475 HHs)	
Parameter	***Static* (1)**	***Dynamic* (2)**	***Static* (3)**	***Dynamic* (4)**
Price $\overline{\beta}$	**0.816** (0.028)	**0.450** (0.037)	**0.810** (0.008)	**0.448** (0.006)
Sugary $\overline{\delta}$	**0.782** (0.016)	**0.492** (0.018)	**0.789** (0.006)	**0.510** (0.003)
Size $\overline{\pi}$	**0.291** (0.021)	**0.175** (0.026)	**0.322** (0.005)	**0.231** (0.004)
Rare $\overline{\theta}$	**-1.039** (0.011)	**-0.500** (0.013)	**-1.050** (0.03)	**-0.536** (0.002)
Folder $\overline{\omega}$	**-0.014** (0.007)	**0.046** (0.010)	**-0.010** (0.002)	**0.035** (0.002)
Rebate $\overline{\varphi}$	**0.079** (0.007)	**0.064** (0.009)	**0.078** (0.002)	**0.067** (0.001)
State Dependence $\overline{\gamma}$	-	**3.033** (0.005)	-	**3.039** (0.001)
Observations	221,554	221,554	9,706,325	9,706,325
Model	Logit	Logit	Logit	Logit
McFadden's R^2	0,09	0,48	0,10	0,48

Notes: *standard errors are in parenthesis. All parameters are significant at the 1% level. Each column displays the parameter estimates of the logit model on different samples, using both the static and the dynamic specification.*

Table 2.A.2: Mixed Logit estimates by the number of Halton sequences

Parameter	**Mixed Logit (200 Hs)**	**Mixed Logit (500 Hs)**	**Mixed Logit (1000 Hs)**
Price $\bar{\beta}$	**2.531**	**2.464**	**2.500**
	(0.053)	(0.057)	(0.058)
Price σ_{β}^{2}	**4.132**	**4.225**	**4.509**
	(0.040)	(0.041)	(0.044)
Sugary $\bar{\delta}$	**0.781**	**0.782**	**0.802**
	(0.030)	(0.030)	(0.030)
Sugary σ_{δ}^{2}	**2.428**	**2.467**	**2.445**
	(0.019)	(0.020)	(0.019)
Size $\bar{\pi}$	**-0.441**	**-0.372**	**-0.342**
	(0.024)	(0.025)	(0.025)
Size σ_{π}^{2}	**2.977**	**2.806**	**2.793**
	(0.019)	(0.017)	(0.017)
Rare $\bar{\theta}$	**-0.973**	**-0.974**	**-0.984**
	(0.013)	(0.013)	(0.014)
Folder $\bar{\omega}$	**-0.026**	**-0.026**	**-0.027**
	(0.011)	(0.011)	(0.011)
Rebate $\bar{\varphi}$	**0.163**	**0.164**	**0.165**
	(0.010)	(0.010)	(0.010)
State dependence $\bar{\gamma}$	**1.854**	**1.839**	**1.834**
	(0.005)	(0.005)	(0.005)
Log-Likelihood	**-246,460**	**-246,060**	**-245,920**

Notes: *standard errors are in parenthesis. All parameters are significant at the 1% level. Each column displays the parameter estimates using 200, 500 and 1000 Halton sequences (Hs).*

Table 2.A.3: Own and cross-price elasticities of demand (Logit)

	A.1	A.2	A.3	B.1	B.2	B.3	C.1	C.2	D.1	D.2	E.1	E.2	F.1	G.1
A.1	**-0.37**	0.04	0.05	0.06	0.04	0.06	0.08	0.06	0.06	0.07	0.08	0.05	0.08	0.07
A.2	0.05	**-0.23**	0.05	0.05	0.05	0.06	0.07	0.06	0.06	0.07	0.07	0.05	0.08	0.07
A.3	0.02	0.01	**-0.30**	0.02	0.02	0.03	0.03	0.02	0.02	0.02	0.03	0.02	0.03	0.02
B.1	0.05	0.03	0.04	**-0.40**	0.04	0.05	0.06	0.05	0.05	0.06	0.07	0.04	0.07	0.06
B.2	0.04	0.05	0.05	0.05	**-0.24**	0.06	0.06	0.06	0.05	0.07	0.07	0.04	0.07	0.07
B.3	0.01	0.01	0.02	0.01	0.01	**-0.33**	0.01	0.01	0.01	0.01	0.00	0.01	0.01	0.01
C.1	0.02	0.01	0.02	0.02	0.01	0.02	**-0.59**	0.02	0.02	0.02	0.03	0.02	0.03	0.02
C.2	0.02	0.02	0.03	0.02	0.02	0.03	0.03	**-0.37**	0.03	0.03	0.04	0.02	0.04	0.03
D.1	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.01	**-0.10**	0.01	0.01	0.01	0.01	0.01
D.2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	**-0.11**	0.00	0.00	0.00	0.00
E.1	0.00	0.00	0.01	0.01	0.00	0.00	0.01	0.01	0.01	0.01	**-0.50**	0.00	0.01	0.01
E.2	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.03	0.03	0.03	**-0.18**	0.04	0.04
F.1	0.01	0.00	0.01	0.01	0.00	0.00	0.01	0.01	0.01	0.01	0.01	0.01	**-0.52**	0.01
G.1	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.01	0.01	0.00	0.01	**-0.29**

Notes: *own-price elasticities are in bold characters. Elasticities are computed at the sample average.*

Table 2.A.4: Parameter estimates of Logit model with observed group heterogeneity

Parameter	Estimate	Demographics	Estimate		
			Price	Sugary	Size
Price $\bar{\beta}$	**14.185** (0.143)	Income	**-1.379** (0.014)	**-0.517** (0.008)	**0.295** (0.015)
Sugary $\bar{\delta}$	**5.280** (0.080)	Education	**-0.148** (0.013)	**-0.094** (0.007)	**-0.048** (0.014)
Size $\bar{\pi}$	**-2.837** (0.151)	Sugar share	**-0.695** (0.046)	**2.257** (0.033)	
Rare $\bar{\theta}$	**-0.547** (0.003)	Household's size	**0.049** (0.002)	**0.062** (0.001)	**0.059** (0.002)
Folder $\bar{\omega}$	**0.029** (0.002)	Kids	**0.018** (0.007)	**0.015** (0.004)	**0.055** (0.008)
Rebate $\bar{\varphi}$	**0.069** (0.002)	Age>70	**0.116** (0.013)	**0.011** (0.007)	**-0.247** (0.014)
State dependence $\bar{\gamma}$	**3.069** (0.001)	70>Age>58	**0.067** (0.011)	**-0.032** (0.005)	**-0.278** (0.011)
		58>Age>46	**0.014** (0.009)	**0.001** (0.005)	**-0.159** (0.009)
		46>Age>34	**0.007** (0.008)	**-0.080** (0.005)	**-0.005** (0.009)
Brand fixed effects			yes		
N° of Observations			4,207,299		
McFadden's R^2			0.50		

Notes: *standard errors are in parenthesis.*

Table 2.A.5: Frequent versus All Customers

Parameter	***All Customers***	***Frequent Customers***
Price $\overline{\beta}$	**2.531**	**2.617**
	(0.053)	(0.004)
Price σ_β	**4.132**	**4.238**
	(0.040)	(0.004)
Sugary $\overline{\delta}$	**0.782**	**0.570**
	(0.030)	(0.024)
Sugary σ_δ	**2.428**	**2.360**
	(0.019)	(0.014)
Size $\overline{\pi}$	**-0.441**	**-0.698**
	(0.024)	(0.019)
Size σ_π	**2.977**	**2.933**
	(0.019)	(0.014)
Rare $\overline{\theta}$	**-0.973**	**-0.990**
	(0.013)	(0.011)
Folder $\overline{\omega}$	**-0.026**	**-0.025**
	(0.011)	(0.008)
Rebate $\overline{\varphi}$	**0.163**	**0.165**
	(0.010)	(0.008)
State Dependence $\overline{\gamma}$	**1.854**	**1.933**
	(0.005)	(0.004)
Model	Dynamic Mixed Logit	Dynamic Mixed Logit
Brand Fixed Effects	yes	yes
Observations	221,554	358,225
McFadden's R^2	0.55	0.57

Notes: *standard errors are in parenthesis. All parameters are significant at the 1% level.*

Table 2.A.6: Robustness check (time-varying brand intercepts)

Parameter	**Dynamic Mixed Logit (1)**	**Dynamic Mixed Logit (2)**
Price $\bar{\beta}$	**2.531** (0.053)	**2.829** (0.053)
Price σ_β	**4.132** (0.040)	**4.048** (0.037)
Sugary $\bar{\delta}$	**0.782** (0.030)	**0.818** (0.030)
Sugary σ_δ	**2.428** (0.019)	**2.436** (0.019)
Size $\bar{\pi}$	**-0.441** (0.024)	**-0.670** (0.027)
Size σ_π	**2.977** (0.019)	**2.832** (0.018)
Rare $\bar{\theta}$	**-0.973** (0.013)	**-0.912** (0.015)
Folder $\bar{\omega}$	**-0.026** (0.011)	**-0.058** (0.012)
Rebate $\bar{\varphi}$	**0.163** (0.010)	**0.183** (0.011)
State Dependence $\bar{\gamma}$	**1.854** (0.005)	**1.844** (0.005)
Time-varying brand effect	no	yes
McFadden's R^2	0.55	0.57

Notes: *standard errors are in parenthesis. All parameters are significant at the 1% level.*

Table 2.A.7: Demand Model with Outside Good (purchase of other sodas)

Parameter	*No Outside Good*	*Outside Good*
Price $\overline{\beta}$	**2.531**	**2.541**
	(0.053)	(0.019)
Price σ_β	**4.132**	**1.746**
	(0.040)	(0.010)
Sugary $\overline{\delta}$	**0.782**	**0.100**
	(0.030)	(0.022)
Sugary σ_δ	**2.428**	**1.982**
	(0.019)	(0.010)
Size $\overline{\pi}$	**-0.441**	**-0.803**
	(0.024)	(0.017)
Size σ_π	**2.977**	**1.278**
	(0.019)	(0.014)
Rare $\overline{\theta}$	**-0.973**	**-0.220**
	(0.013)	(0.011)
Folder $\overline{\omega}$	**-0.026**	**-0.042**
	(0.011)	(0.009)
Rebate $\overline{\varphi}$	**0.163**	**0.429**
	(0.010)	(0.008)
State Dependence $\overline{\gamma}$	**1.854**	**2.271**
	(0.005)	(0.004)
Model	Dynamic Mixed Logit	Dynamic Mixed Logit
Brand Fixed Effects	yes	yes
Observations	221,554	333,097
Share Outside Good	NA	0.33
McFadden's R^2	0.55	0.58

Notes: *standard errors are in parenthesis. All parameters are significant at the 1% level.*

2.3.A Demand Curves (long-run)

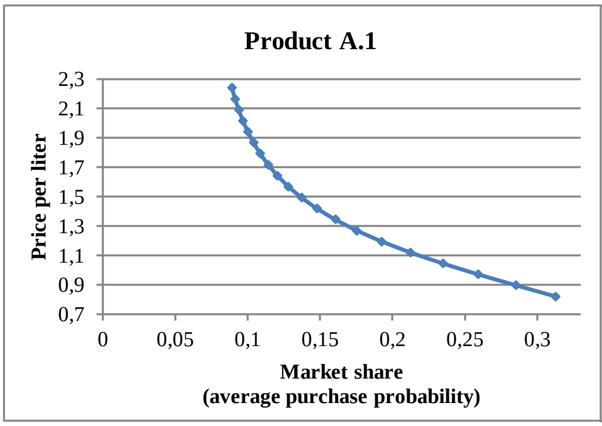

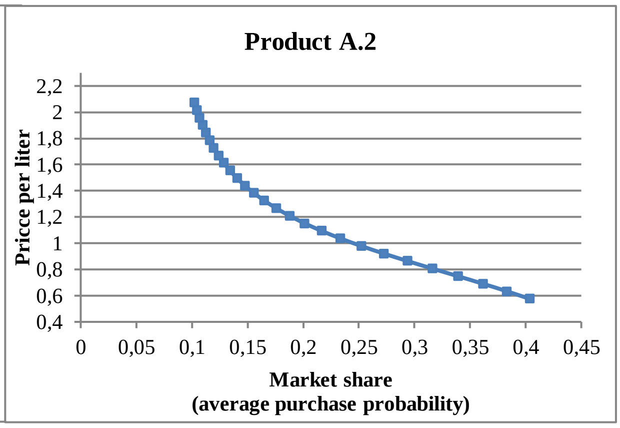

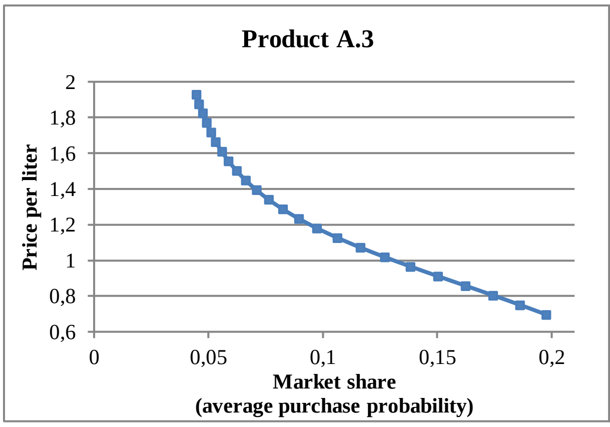

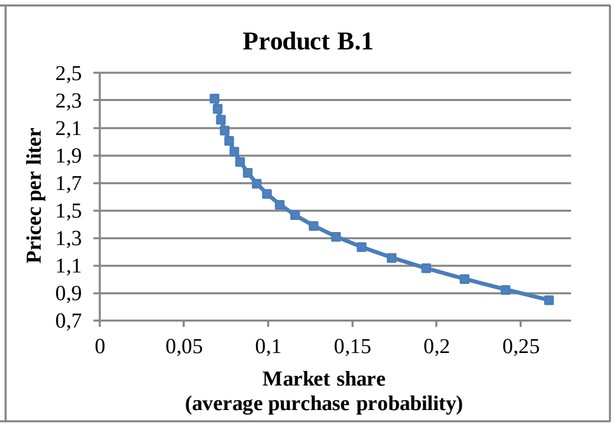

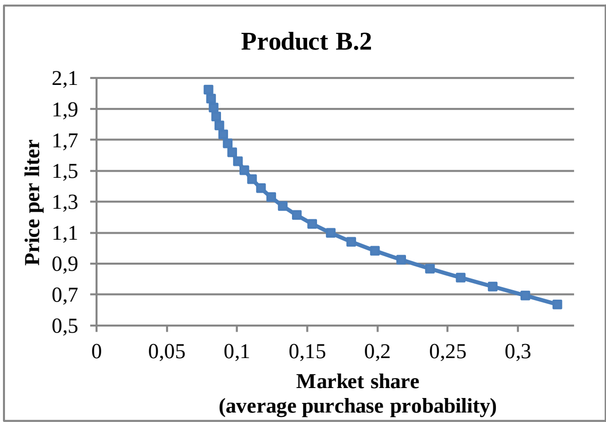

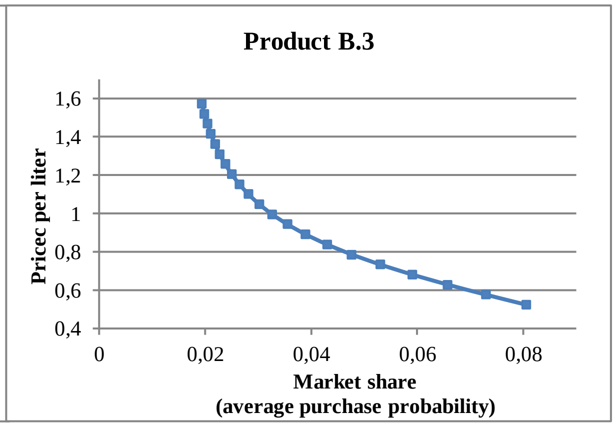

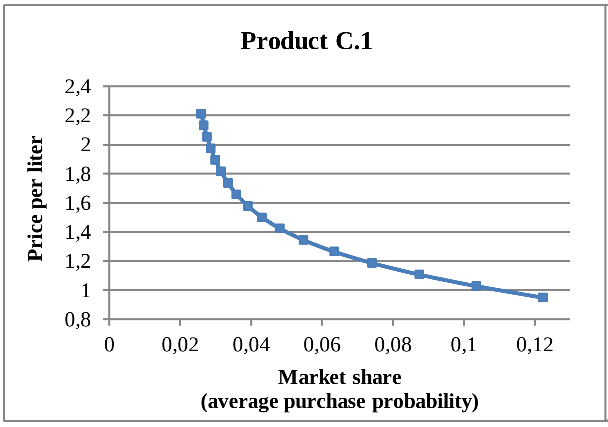
Product C.1
2,4
2,2
2
1,8
1,6
1,4
1,2
1
0,8
Price per liter
0
0,02
0,04
0,06
0,08
0,1
0,12
Market share
(average purchase probability)

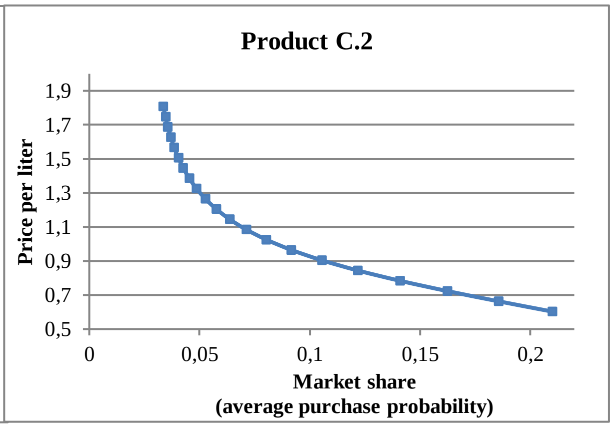
Product C.2
1,9
1,7
1,5
1,3
1,1
0,9
0,7
0,5
Price per liter
0
0,05
0,1
0,15
0,2
Market share
(average purchase probability)

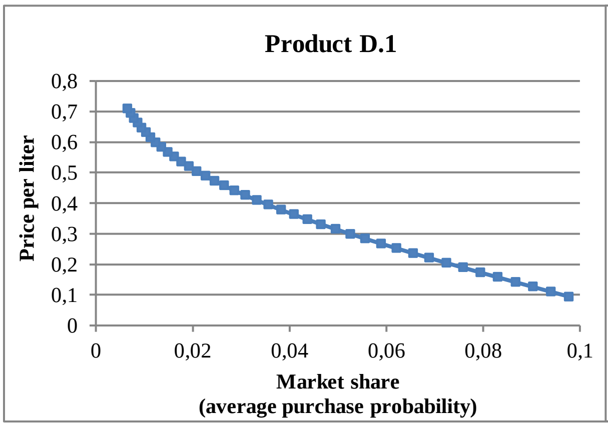
Product D.1
0,8
0,7
0,6
0,5
0,4
0,3
0,2
0,1
0
Price per liter
0
0,02
0,04
0,06
0,08
0,1
Market share
(average purchase probability)

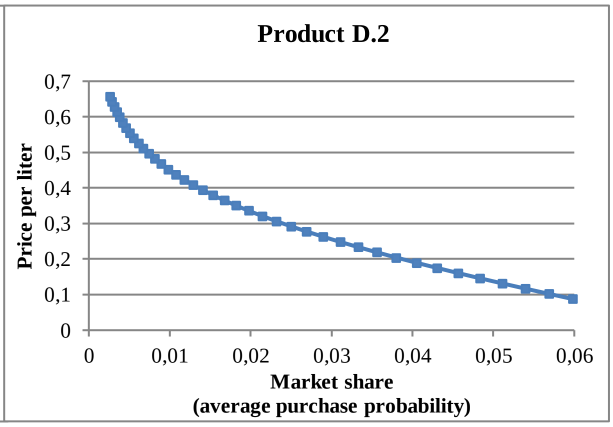
Product D.2
0,7
0,6
0,5
0,4
0,3
0,2
0,1
0
Price per liter
0
0,01
0,02
0,03
0,04
0,05
0,06
Market share
(average purchase probability)

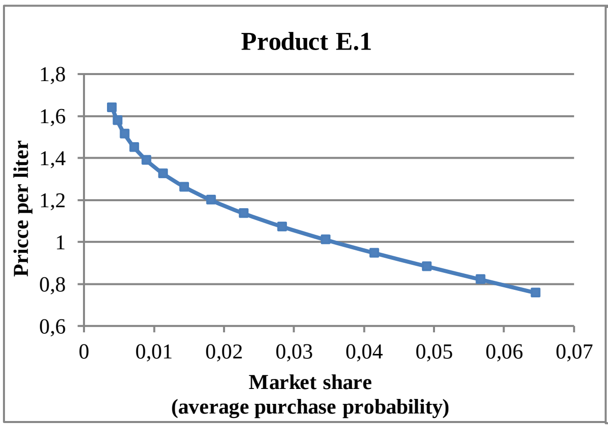
Product E.1
1,8
1,6
1,4
1,2
1
0,8
0,6
Pricce per liter
0
0,01
0,02
0,03
0,04
0,05
0,06
0,07
Market share
(average purchase probability)

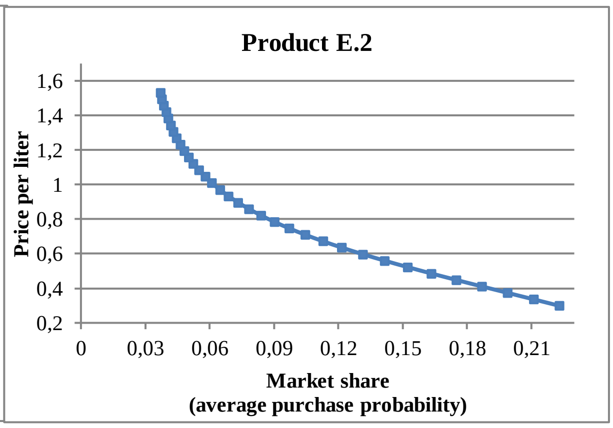
Product E.2
1,6
1,4
1,2
1
0,8
0,6
0,4
0,2
Price per liter
0
0,03
0,06
0,09
0,12
0,15
0,18
0,21
Market share
(average purchase probability)

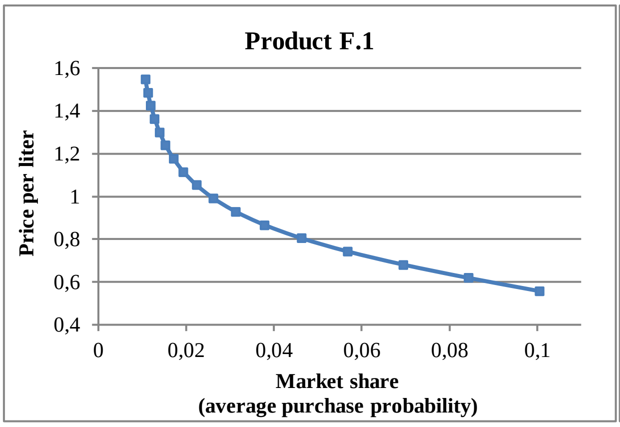
Product F.1
1,6
1,4
1,2
1
0,8
0,6
0,4
Price per liter
0
0,02
0,04
0,06
0,08
0,1
Market share
(average purchase probability)

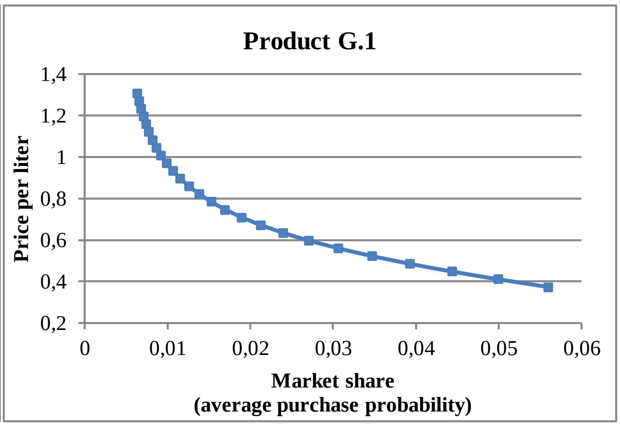
Product G.1
1,4
1,2
1
0,8
0,6
0,4
0,2
Price per liter
0
0,01
0,02
0,03
0,04
0,05
0,06
Market share
(average purchase probability)

Chapter III

Heterogeneity in the Tax Pass-Through to Spirit Retail Prices: Evidence from Belgium*

Authors: Jean Hindriks† and Valerio Serse‡

Abstract: On 1st November 2015, the Belgian government increased the excise tax on alcoholic beverages. For spirits with 40% of alcohol and bottle size of 70cl, this tax change is equivalent to an amount of 2,43€ per bottle of spirits. This paper studies the impact of this tax reform at the store level on the (posted) retail price of six major brands of spirits, using a difference-in-differences method. The estimation is based on *a balanced panel of scanner data* from a major supermarket chain (with a 33% market share) and uses the retail prices of the same brands sold in France by the same supermarket chain as a control group. Having information on each store location, we show spatial variations in the tax pass-through for *homogenous* products. We find that these variations are strongly related to the intensity of local competition and to a lesser extent to the proximity to the borders (mainly with Luxembourg which is the low-price country). We find that the tax was quickly passed through during the first month of tax implementation and that it was mostly over-shifted. However, we also find that both the border and the competition effects are not instantaneous but arise several months after the tax reform. These findings have important implications for alcohol control policies as they highlight that the incidence of alcohol taxation can vary greatly across space and affect differently households depending on where they live.

JEL No: H2, H22, H32, H71, I18.

Keywords: tax pass-through, scanner data, competition, cross border shopping.

* We have benefited from very useful comments and suggestions of two anonymous referees and the Editor Camille Landais. We are also grateful to Andreu Arenas, Liran Einav, Rachel Griffith, Christian Hafner, Johannes Johnen, William Parienté, Isabelle Thomas and seminar or conference participants in various places for helpful discussions. The financial support from the FNRS is gratefully acknowledged. This version of the paper was also published in the *Journal of Public Economics*, *176*, 142-160.

† CORE, Université catholique de Louvain, e-mail: jean.hindriks@uclouvain.be.

‡ CORE, Université catholique de Louvain, FNRS-FRESH grant holder, e-mail: valerio.serse@uclouvain.be.

3.1 Introduction

On 1st November 2015, the Belgian government increased the excise tax on alcoholic beverages. This tax reform was not in reaction to market conditions (providing an exogenous change in tax rates). It was part of the general governmental tax shift plan aiming to shift the tax burden from labor to consumption (with higher taxes on electricity, gasoil, cigarettes, alcoholic beverages, and sodas). The tax increase was different across alcohol types. For instance, the taxes on beer and wine have increased by 8.5% and 31%, respectively. The strongest tax increase was for the category of spirits, which is also the category that was taxed most heavily before the tax reform. From 2.127,68 €/hl per % alcohol to 2.992,79 €/hl per % alcohol.[38] That is, an increase of 41% in excise tax. Considering a standard bottle of 70 cl with 40°C, this tax change amounts to an extra tax of 2,43€ per bottle.[39] This tax reform was heavily criticized in the media for inducing sales loss and its failure to bring extra revenues. In fact, the total revenue from excise taxes on alcohol was 318€ million in 2015, 323€ million in 2016 and 319€ million in 2017.[40] One survey of 425 local retailers organized during the spring of 2016 by the SNI (Syndicat neutre des indépendants - Trade union for independent workers) suggests that sales have declined by 14 % in volume and shop thefts have increased by 11% in a single year. The federation of spirit and wine (Vinum & Spiritus) blames cross-border shopping for the loss of sales, since the tax reform has considerably increased the relative price of Belgian spirits with respect to all the neighboring countries. For example, the price of a bottle of Gin Gordon was after the reform 15 euros in Belgium against 9 euros in Luxembourg. Given that 50% of the Belgian households live within a distance of 50 km from the border we could indeed expect massive cross-border shopping.

[38] In comparison with neighbouring countries, excise taxes (and VAT tax levied on the price inclusive of the excise duty) are as follows: Belgium 2.992 €/hl (VAT 21%), France 1.741 €/hl (VAT 20%), The Netherlands 1.686 €/hl (VAT 21%), Germany 1.303 €/hl (VAT 19%), Luxembourg 1.041 €/hl (VAT 17%). (European Commission 2018, Excise duty on alcohol beverages).

[39] The magnitude of the tax change on beer and wine is much lower than on alcohol. The tax change for a standard bottle of wine (75cl) is 0,13€, and for a can of beer (33cl) it is 0,01€. Interestingly, such differentiation of the tax changes is consistent with Griffith et al (2019) who estimates the welfare gains from varying tax rates across different types of alcohol depending on the concentration of alcohol externalities among heavy drinkers. It is interesting to concentrate on the spirit market because the planner can target the most socially harmful drinking by taxing more heavily the ethanol in products that are disproportionally consumed by problem drinkers (see Griffith, 2019).

[40] Source: SPF Finances Belgium. Available at <https://finances.belgium.be/fr/statistiques_et_analyses/rapport-annuel/chiffres/2018/budget-recettes/recettes-ag-douanes-et-1>

The magnitude of this tax increase provides a unique opportunity to estimate the tax-pass-through of spirits in the Belgian market and to focus on spatial heterogeneity in tax incidence across geographical areas. Understanding the incidence of alcohol taxation is fundamental to assess the effectiveness of this policy to improve public health and/or generate fiscal revenues. This is also important to identify how the tax burden and health benefits are distributed in the population. Although alcohol is typically taxed homogeneously within a given jurisdiction, the extent to which a tax is passed through to alcohol retail prices can be substantially heterogeneous across geographical areas. Theoretically, the tax pass-through is a function of both the price elasticity of demand and the structure of the supply-side of the market. Spatial differences in these two factors can therefore explain the heterogeneity in tax incidence within a tax jurisdiction. The proximity to a lower taxed state can be another important determinant of tax shifting due to tax avoidance by means of cross-border shopping.

This paper contributes to the empirical literature on tax pass-through by analyzing the impact of the recent alcohol tax reform in Belgium on spirit retail prices using *a balanced panel of supermarket scanner data* from a major group of retailers. Unlike conventional scanner average price data used in the literature (e.g. Nielsen measured prices), we use more detailed data on posted prices from a large retail chain. Posted prices are the same as actual transaction prices in general. However, they may differ relative to transaction prices because they are not conditional on purchase and thus less sensitive to local and cyclical shocks (Coibion et al., 2015). Posted prices are not dependent neither on measurement errors due to loyalty cards (Einav et al., 2010). That does not mean that posted prices are always superior to actual transaction prices. Indeed posted prices could in principle not be updated, and this would not be realized if there are no transactions. In our case, posted prices are automatically updated daily for each item. Although posted price data are only observed for all the retailers of the same supermarket chain, this group possesses a significant market share (about one third) and is publicly committed to match prices of local competitors (price matching strategy). Hence, their price can be considered as representative of the general price evolution in the market.[41] Furthermore, as this group is also present in other countries, price data for the exact same products in France (not submitted to the tax change) can be used as a control group. This allows

[41] The local competition analysis would have been more interesting if we could observe the price changes in other stores, to see if the price changes are coordinated. In the same vein, the cross-border shopping would also have benefited from the possibility to observe the price changes in stores just next but on the other side of the border.

measuring the tax pass-through to spirit retail prices by means of a *“difference-in-differences”* estimator.

In our analysis, posted prices include any taxes. This is different from several studies in the U.S. where posted prices do not include some taxes. The question that arises is about the salience of the tax change. Tax salience matters to assess the tax pass-through since we may expect the retailer to shift more of the tax when consumers do not know whether the tax change has occurred (the tax is less salient). Chetty et al. (2009) provide experimental evidence that consumers are less sensitive to (non-posted) tax changes than they are to changes in the posted price. Interestingly, in our study the tax change was explicitly announced, and the posted prices include the tax. Thus, we may expect the tax to be more salient. Nevertheless, our results suggest significant tax over-shifting.

The rich nature of the dataset allows testing for and explaining spatial heterogeneity in tax pass-through over Belgium. Having information on both proximity to the border and the number of competitors for each store, this work provides new evidence on the effects of cross-border shopping and the intensity of competition on the pass-through of alcohol excise taxes. Yet, we cannot make causality claim here because we do not have exogenous variations in the intensity of local competition at the retail level during the tax reform.[42] As price data are collected over several months, this study also checks for the evolution of the tax pass-through over each month after the tax hike and tests whether the observed heterogeneity in price hikes is permanent or just temporary.

The spatial dispersion in posted prices and in the tax pass-through contrasts with the recent empirical study on uniform pricing in U.S. retail chains based on the Nielsen price measure (see Della Vigna & Gentzkow, 2017). The difference may result from the uniform mark-up rule regulation used in the U.S. (Miravete et al., 2017). These findings highlight that the incidence of alcohol taxation can vary greatly across geographical areas, even within a small country as Belgium. We find that the stores’ heterogeneity in tax shifting is strongly related to local differences in the intensity of competition at the retail level. Surprisingly enough, we do not find that differences in tax shifting are significantly related to the proximity to the border in general. Although the tax reform has considerably increased the relative price of Belgian spirits with respect to all the neighboring countries, we find a lower tax shifting only in stores bordering Luxembourg. Which is the neighboring country with lowest spirit prices before the

[42] This would have allowed us to purge the effect of competition from unobserved differences across stores that are specific to their location.

alcohol tax reform. In line with the previous literature, we find that the tax was quickly shifted to spirit retail prices. With a significant tax over-shifting already during the first month of tax hike. Interestingly, we find that both the border and the competition effects are "back-loaded" in the sense that they show up with some lag (few months after the reform). This suggests that it took some time before stores adjusted their prices to both the foreign and domestic competitors.

The rest of the paper is organized as follows. In Section 3.2, we provide a review of the relevant empirical literature focusing on tax pass-through and identify our contribution. In Section 3.3, we provide a brief account of the theory on the tax pass-through and how it relates to market structure and the shape of the demand. In Sections 3.4 and 3.5, we describe our dataset and perform the empirical analysis. Section 3.6 provides some summary statistics about the demand response (change in the quantity of bottles sold) to the tax hike. Section 3.7 concludes.

3.2 Contribution to the literature

Various empirical studies estimate the tax pass-through to the retail price of sin goods. In particular, recent works focused on tax pass-through in the market of sodas (Cawley & Frisvold, 2015; Berardi et al., 2016; Campos-Vazquez & Medina-Cortina, 2016; Grogger, 2017), cigarettes (Harding, Leibtag & Lovenheim, 2012; DeCicca, Kenkel & Liu, 2013; Xu et al., 2014) and alcoholic beverages (Kenkel, 2005; Carbonnier, 2013; Ally et al., 2014; Conlon & Rao, 2016; Shrestha & Markowitz, 2016). These studies mostly consist of reduced-form analysis that use price data collected from different sources during a period of tax policy change. The common strategy is to regress the price variable on a tax indicator plus a set of controls in order to isolate the causal impact of the tax on prices. [43]

Part of this literature, however, identifies tax pass-through by means of a *"difference"* estimator (see DeCicca, Kenkel & Liu, 2013; Xu et al., 2014; Kenkel, 2005; Carbonnier, 2013; Ally et al., 2014; Conlon & Rao, 2016). That is, by measuring pre- versus post-tax difference in retail prices. Some of the most recent papers overcome this limitation by introducing control groups that account for the counterfactual price evolution in absence of tax policy change. This allows estimating the tax pass-through by means of a typical *"difference-in-differences"* estimator. Nevertheless, type and quality of control groups for prices tend to vary over different studies.

[43] Sources of price data can include, for instance, online price comparison services (Ally et al., 2014; Berardi et al., 2016), self-reported purchases (DeCicca, Kenkel & Liu, 2013; Xu et al., 2014), scanner data (Harding, Leibtag & Lovenheim, 2012; Conlon & Rao, 2016), governmental agencies (Campos-Vazquez & Medina-Cortina, 2016; Grogger, 2017) and telephone interviews (Kenkel, 2005).

For instance, Berardi et al. (2016), which estimates the impact of the “soda tax” on prices in France, use the price of untaxed beverages as a control group for the taxed products. The same approach is adopted by Campos-Vazquez & Medina-Cortina (2016) and Grogger (2017), which both study the pass-through of the “soda tax” implemented in Mexico in January 2014. Conversely, Harding, Leibtag & Lovenheim (2012), who analyze the pass-through of cigarette excise taxes in the United States, use as a control group the same cigarette products sold in those states that did not change their cigarette excise taxes. Similarly, Cawley & Frisvold (2017) use as a control group the price of sugar-sweetened-beverages (SSBs) in the city of San Francisco to estimate the pass-through of the tax on SSBs implemented in the neighboring city of Berkley, California.

This literature generally finds that tax incidence is quite heterogeneous across products and that all three patterns of under-, over- and full shifting are likely to occur after the implementation of a tax on sin goods. In the context of alcohol taxation, existing evidence generally suggests tax over-shifting with a large heterogeneity of tax pass-through across products. Kenkel (2005) find that the pass-through of the alcohol tax hike occurred in Alaska in 2002 ranged between 167% and 213% for 6 major brands of distilled spirit. Ally et al. (2014) estimate the pass-through of excise duties and VAT in UK during the period 2008-2011. They find evidence of tax over-shifting for spirits on average, but they also find a significant tax under-shifting for the cheapest brands. This evidence highlights the complexity in designing sin taxes aimed at improving public health. As price hikes tend to differ even within the same category of taxed products, there should be a rising concern about both the substitution effect towards other taxed goods and the distribution of tax incidence across different types of consumers. Our paper extends this literature by providing evidence of a further dimension of heterogeneity in alcohol tax shifting. That is, the spatial heterogeneity in the tax pass-through for *homogeneous* products. Although such heterogeneity in tax shifting can be theoretically explained by differences in price elasticities and market structure across geographical areas (Hindriks & Myles, 2013), little attention has been given to this phenomenon in the empirical literature. In this paper, we focus on two possible determinants of spatial heterogeneity in tax shifting: the variation in the scope for cross-border shopping and the variation in the local intensity of competition at the retail level.

Prior empirical papers on cross-border shopping have studied the demand side. That is how price differences create incentive to cross the border line (see, for instance, Gopinath et al., 2011; Asplund et al., 2007; Manuszak & Moul, 2009; Chandra et al., 2014 and Chiou &

Muehlegger, 2008). This empirical work has shown that consumers do respond to price differences by engaging in cross-border shopping. What is less studied is how retailers in turn respond to that cross-border shopping. Harding, Leibtag & Lovenheim (2012) and Cawley & Frisvold (2017), use price data at the store level, respectively for cigarettes and sodas, to find that part of tax pass-through heterogeneity across stores can be explained by their proximity to states with lower tax rates on cigarettes and sodas. In particular, they find lower tax pass-through in stores next to the border, thus suggesting that the scope for cross-border shopping drives down the extent to which stores can rise prices after a tax hike. Doyle & Samphantharak (2008) study the effects of cross-border competition on the gasoline tax shifting to retail prices. They use data of daily prices at the gas station level to estimate the impact of a temporary suspension, and a subsequent reinstatement, of the gasoline sales tax in Illinois and Indiana on the retail price of gasoline, which followed a price spike in the spring of 2000. They adopt a difference-in-differences approach by using the gasoline retail price of neighboring states as control group. Their findings on the border effect are mixed but overall they suggest a smaller tax shift for gas stations close to the border, especially for the reinstatements (tax increase), with some evidence of tax spillover across state borders.

Like these studies, the contribution of our paper is on the cross-border shopping effect on prices. We study how the distance to the border affects the extent of the tax shifting to spirit retail prices. Understanding the tax shifting for alcoholic beverages at the border provides precious insights into how tax avoidance can reduce the effectiveness of the sin tax in curbing the consumption of alcohol or generating tax revenues. Most papers analyzing the effectiveness of alcohol taxes to curb demand get results on volume sales that are only valid conditional on the tax incidence on prices (Wagenaar et al, 2008). With cross-border shopping, affected stores might be less willing to pass on the tax in order to avoid losing consumers to nearby (untaxed) stores. Belgium is a nice candidate for this analysis because it is a small country with high population density and a sizeable population at a short distance to the borders with four neighboring countries using the same currency (Euros). Unlike the previous literature, we also study the timing of the border effect on the tax pass-through. We show that this has to be carefully taken into account in empirical works as it may take time for stores to internalise the cross-border shopping in their price adjustment to the tax reform.

It is important to mention that in this paper, we do not estimate the cross-border spillover effect of the tax change in the neighboring stores on the other side of the border. Bajo-Buenestado & Borrella-Mas (2018) provide interesting estimates (using differences-in differences) of this "*cross-border pass-through*" from the fuel tax reform in Portugal on the Spanish fuel prices of stations that are close to the Portuguese-Spanish border. Their control group are the Spanish gas stations that are far from the border.[44] In our paper, we only consider the "*domestic pass-through"* since we do not observe price changes in foreign stores near the border (our control group are French stores that are far from the border).

In this paper, we also study how variation in competition at the store level may relate to the spatial variations in tax shifting. Economic theory indicates that the intensity of competition can extensively affect the extent of tax pass-through to retail prices. Yet, this competition effect is not very much studied in the empirical literature. Doyle & Samphantharak (2008) estimate how the tax shifting to gasoline retail prices varies across local markets with different levels of brand concentration. The idea is that the tax change should be reflected upstream in the wholesale price depending on the market power in the wholesale market. They measure the share of gas stations for each (wholesale) brand in a local market and compute a Herfindahl–Hirschman Index of brand concentration. They find some evidence that tax shifting varies with brand concentration at the ZIP code level, with the price hike (after the tax reinstatement) being 2 percentage point lower in the least concentrated markets. Stolper (2016) finds that tax pass-through differences range from 70% to 120% at the specific station level in the Spanish fuel market. Greater market power measured by brand concentration is strongly associated with higher pass-through, even after conditioning on detailed demand-side characteristics.

Campos-Vazquez & Medina-Cortina (2016), using price data at the store level, show that the competitive barriers faced by each store generate significant differences in the shifting of the "soda tax" in Mexico. They use as control group the water price that is not subject to the tax increase, but whose price is highly correlated with prices of the taxed product, the soft drinks (treated group). They compute the number of competing retailers within a distance of 8km from each store and find that the tax pass-through decreases with the number of competitors. Etilé et al (2018) find similar result for the 2012 soda tax in France. We extend this literature by providing evidence of the competition effect on the tax shifting to spirit prices using as a control group the same product sold by the same chain in a different country not subject to the tax hike. Although Belgium is a relatively small country, we find a very large store heterogeneity in tax pass-through that can be related to differences in competition intensity at the retail level. We

[44] Doyle & Samphantharak (2008) do a similar analysis for the US and provide evidence of cross-border pass-through.

also provide novel evidence about the timing of this effect and show that the competition effect is back-loaded and arises with some lag.

Lastly, evidence on the tax pass-through timing suggests that prices tend to react quickly to the introduction of excise taxes. The “soda tax” in Mexico in January 2014 was already fully shifted into soda prices during the first month of implementation (Campos-Vazquez & Medina-Cortina, 2015; Grogger, 2017). While the “soda tax” in France in January 2012 was gradually passed through to retail prices and fully shifted after six months (Berardi et al., 2016). Carbonnier (2013) reports that the increase in excise taxes on alcohol implemented in France in January 1997 was immediately fully shifted to the price of both beer and aperitif during the first month of tax hike. Conlon & Rao (2016) find that excise taxes on distilled spirits in the U.S are shifted within a month and are often over-shifted. Our paper confirms those findings of a quick tax shifting with frequent over-shifting.

3.3 Theoretical Framework

The basic theory on tax incidence in industrial organization is about estimating the changes in prices and profits resulting from a tax (Fullerton & Metcalf, 2002). Let us denote the excise tax t and the producer price $p(t)$, then the consumer price is $q(t) = p(t) + t$. In our context of supermarket transactions, the producer should be understood as the retailer. Under perfect competition, the tax incidence is very simple. The tax shifts the supply curve vertically upward by the amount of the tax. The incidence of the tax on prices is $q'(t) = p'(t) + 1$ where $q'(t)$ and $p'(t)$ are the tax derivative of the consumer and producer prices. The extent to which consumer price rises is determined by the elasticities of the supply and demand curves. Formally, the pass-through rate is given by:

$$dq/dt = \frac{1}{1 + \left(\frac{\varepsilon_D}{\varepsilon_S}\right)}$$

where ε_D is the elasticity of demand (in absolute value) and ε_S is the elasticity of supply (Weyl & Fabinger, 2013). If the demand is inelastic, $q'(t) = 1$ and thus $p'(t) = 0$, that is consumer price will rise by the exact amount of the tax and producer price is unchanged. We have perfect tax shifting. In all other cases the consumer price increases to a lesser extent than the amount of the tax $q'(t) < 1$, and the producer price decreases $p'(t) < 0$. The tax is shifted in part to the consumer and in part to the producer as a function of the elasticities of supply and demand. In this general case we have tax under-shifting $q'(t) < 1$. Hence, with perfect competition, the full amount of the tax may be shifted to consumers but never more, and this is only possible if the demand is perfectly inelastic.

Under imperfect competition, tax incidence is different and tax over-shifting becomes possible. This possibility depends on the shape of the demand function. To illustrate that point we need to trace the effect of the tax on the profit-maximization decisions of the imperfectly competitive firms (here retailers). To see that easily, we follow Hindriks & Myles (2013). Consider a monopoly situation with constant marginal cost. Figure 3.1a depicts the profit maximization of a monopoly choosing not shifting all the tax on the consumer. Indeed, the tax is shown to move the intersection between marginal cost and marginal revenue (i.e. the profit maximization condition) from a to b with a reduction of output from $y°$ to y^t and consumer price rises from p to q. In this case, price rises by less than the tax imposed $(q - p < t)$.

Figure 3.1a: Tax Under-Shifting

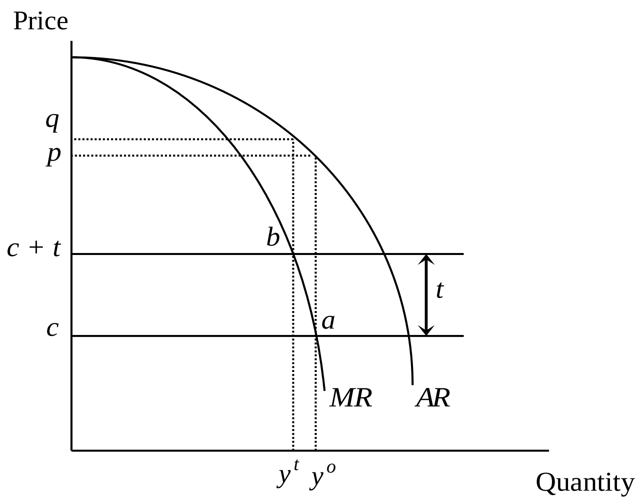

Figure 3.1b: Tax Over-Shifting

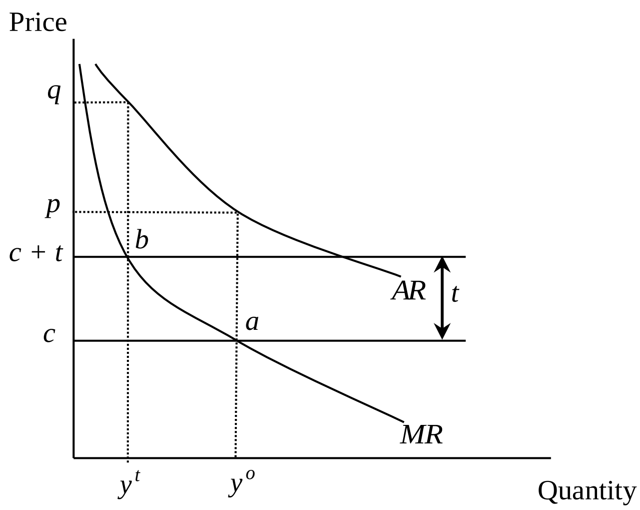

In contrast, Figure 3.1b depicts the same monopoly facing a demand function with a different shape. The demand has a convex shape: it becomes increasingly flat as quantity increases (whereas, in Figure 3.1a the demand has a concave shape: it becomes increasingly steep as quantity increases). In this case, the tax induces a price increase from p to q that is greater than the amount of the tax $(q - p > t)$, so we have tax over-shifting.

To extent this result to the case of imperfect competition (Cournot-oligopoly), we can consider an isoelastic demand function $X = q^{\varepsilon}$ where $\varepsilon < 0$ is the price elasticity of demand. With a constant price elasticity, the mark up is constant

$$\mu^0(n) = \frac{n}{n - \left(\frac{1}{|\varepsilon|}\right)}$$

where n is the number of (equal-size) competing firms. When firms have different market shares ($s_i > 0$) we replace the number n by n* (with n*<n) the equal-size equivalent Herfindahl index (with $H(n) = \sum_{i=1}^{n} s_i{}^2 = \frac{1}{n*}$). Since $|\varepsilon| > 0$, we have $\mu^0 > 1$. The equilibrium price is obtained by applying the mark up to the marginal cost-plus tax, to get $q(t) = \mu^0(n)[c + t]$. The tax incidence on price is then $q'(t) = \mu^0 > 1$. Hence, there is always tax over-shifting with isoelastic demand and imperfect competition. This is true for $n = 1$ (monopoly) and $n > 1$(oligopoly). In addition, from the expression for the markup, we have that $\mu^0(n)$ is decreasing in n, so as the intensity of competition increases (n increases) the markup decreases reducing the extent of over-shifting. At the limit as competition becomes more and more intense $\mu^0(n)$ tends to 1 and the competitive outcome of perfect tax shifting arises $q'(t) = 1$.[45] Given this markup formulae we expect stores facing more competition and stores facing more elastic demand (cross-border shopping) to shift less of the tax on the retail price.

On the effect of cross-border shopping we would expect that the shifting of the tax to the consumers will be lesser the greater the scope for cross-border shopping into another jurisdiction with unchanged tax. Bajo-Buenestado & Borrella-Mas (2018) propose a theoretical model with imperfect competition among differentiated products and cross-border tax spillover to predict that proximity to the border (interpreted as a reduction in product differentiation) reduces the tax-pass through.

3.4 The Data

The data used in this work are provided by a major Belgian supermarket chain with a market share of 33% in Belgium. This retail chain controls more than 400 local retailers in Belgium, France and Luxembourg. Posted price data are automatically collected by the retailer on a daily basis for every item sold in each store of the group, together with information about any price promotions and rebates. Posted prices differ from the average "measured" price commonly available in scanner data (e.g. Standard Nielsen scanning data price measure in the US). The average "measured" price in a given week is the weekly ratio of sales revenue to the quantity sold. It is a quantity weighted average of posted prices. It can vary across stores and location

[45] The use of price rather than quantity as a strategic variable (Bertrand competition) intensifies competition and reduces profits. This means that the effective elasticity of demand is likely to be larger in magnitude than in the Cournot competition. However, if the cross-price elasticity is limited, the substitutability is limited (differentiated products) then the Cournot markup rule is likely to work. It is also likely to work in markets where competition is stable with no dynamic price wars in general. This kind of stable pricing would arise if firms have been competing for a long time and if there is some kind of price matching strategy in place. Recall that in our case, the supermarket chain under consideration is using an explicit price matching strategy based on local competition.

even though the posted price is uniform. Indeed, stores facing less elastic demand (or higher income) would sell a relatively larger share at higher price, which induces a higher weight on higher prices and thus a higher average price in those stores (see Della Vigna & Gentzkow, 2017).

As stores are located in different areas, posted prices tend to vary considerably both within and across countries. Interestingly, this retail chain acts as local price followers: it is publicly committed to constantly monitor competitors' prices and sell its products at the lowest price in all its local markets. The price monitoring is done either online or manually by a team of its employees. This monitoring occurs on a daily basis and it is on two different levels. First, the chain monitors the price of large retailers operating on a national scale and that price uniformly over the country. If it finds that any of these retailers has a lower price for any of its products, then the chain updates immediately its price on a national scale by setting it just below the price of the competitor. Second, for all remaining small local retailers and the stores not adopting uniform pricing nationally, the chain monitors the local prices of those retailers close to each of its own stores. If the price of a product is lower in any of these local retailers, the chain updates the price of its local store by setting it just below the price of its local competitor.

To inform its customers about the effectiveness of this pricing strategy, any time the price of an item has been recently decreased (less than a week) to match the lowest price of a national/local competitor, the store signals the price change on the price tag and displays the new price in red color. Furthermore, the company regularly publishes on its website a prospectus indicating the average price difference between its stores and the main competitors for every geographical area. The effectiveness of this strategy is also confirmed by the independent Belgian consumer association *Test-Achat*, which provides yearly comparative price reports of Belgian retail chains. For every year that the survey was carried out, they find that our retail chain was the cheapest among all its major competitors.

Given this local price matching strategy, any observed price change in these stores actually reflects a change in the lowest price offered by other retailers in their local market. This allows us to extend the study of the tax pass-through from one specific retail chain to each local market, by including the local influence of other retailers. Although this retail chain is committed to match local prices, one concern that can arise is that local competition in prices is meaningless if the majority of the market is supplied by large retail chains that price nationally. However, this does not seem to be the case for Belgium, where still a large part of grocery stores is made up of small independent local retailers that do not belong to larger retail chains. According to a recent report by Nielsen (2017), the market structure of grocery stores in Belgium is as follows: 42% of stores are small traditional local shops; 37% of stores belong to large and medium sized retail chains, and 21% of stores are other small supermarkets operating either at the local or national level. This market structure suggests some possible degree of spatial price dispersion and opens up to the possibility of heterogeneity in tax incidence over the country.

This work focuses on assessing the tax incidence of the tax hike in Belgium on spirits retail prices by selecting six major brands of spirit that have the unique characteristic of being sold both in Belgium (in 337 stores) and in France (in 71 stores of the same supermarket chain). This allows performing a difference-in-differences analysis by considering the price evolution of the same brand sold in France as control group during the period of tax implementation. We therefore assume that, had the tax not been implemented, the Belgian price of each of these products would have followed the same trend as that one of the same product in France. French prices in the same supermarket chain can be considered as a good control group given that these products share the same cost components and are sold by the same retailer in these two neighbouring countries. Figure 3.A.1 in the appendix shows the location of control stores in France. As French stores are located far away from the Belgian border, we should not expect the Belgian tax reform to impact French prices via cross-border shopping. The French store closest to Belgium is about 70 km away from the Belgian border. Cross-border shopping is unlikely because of both a long driving distance (around one hour) and the fact that French stores in this area (Lorraine region) are much closer to Luxemburg, which is the relevant cross-border shopping destination given its lower spirit prices.

We restrict attention to three brands of vodka, one brand of whiskey and two brands of rum. These products are among the leading brands in the market of spirits and have the unique characteristic of being sold in the same format in many stores both in Belgium and in France. This provides the opportunity to compare the price evolution of the exact same product in these two countries. These products differ in their alcohol content, being either 40% or 37,5% and all products considered have the same bottle size of 70cl. Hence, the tax change is different across these products. For spirits with the 40% of alcohol content the tax increase amounts to 2,43€ per bottle. While for those with the 37,5% this amounts to 2,28€. As mentioned in the introduction, the tax change on spirits was not in reaction to some pre-existing market conditions, as it was part of a general plan of the Belgian government aiming at shifting the tax burden from labor to consumption. This provides us with an exogenous tax reform.

Figure 3.2: Evolution of spirit prices

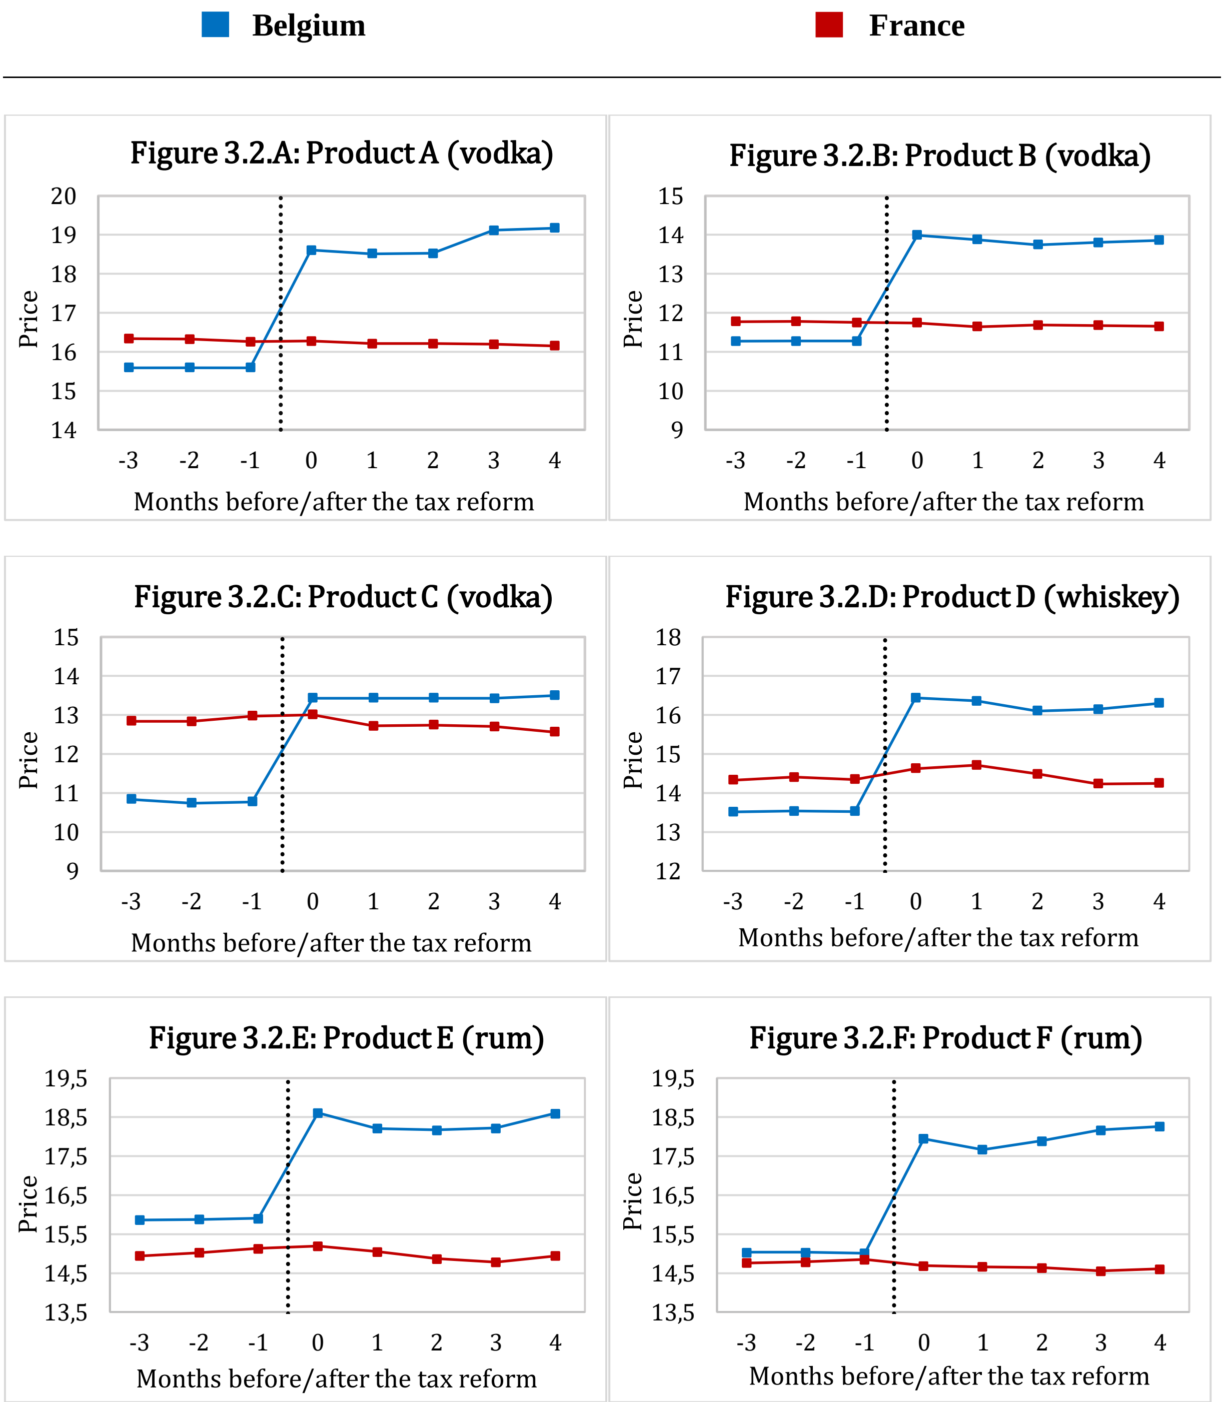

The price data consists of the monthly posted price of each brand of spirit sold in every local store net of any rebate and temporary price promotion. For most products, these discounts are quite frequent during Christmas period, but can also occur in other periods of the year. To control for temporary price promotions, we use the highest daily price of the month (peak price) for each store. This allows controlling for temporary price cuts that are not relevant for the estimation of the tax pass-through to spirit prices.[46] Price records begin three months before the tax reform and end five months after. Figures 3.2.A to 3.2.F display the evolution of the monthly price for each spirit during this period for both French and Belgian stores. Although a longer price series would be preferred to check for common pre-treatment trend, these figures show that prices in both countries did not diverge over the 3 months prior the tax hike. This gives us a first check of the validity of the control group. As it can be seen from these figures, the tax reform impacted Belgian prices immediately the month of its implementation, while French prices stayed quite stable all over the period. Interestingly, for products A to D, the tax reform reversed the price differential between French and Belgian stores. Those products were cheaper in Belgium before the reform and became more expensive after.

Table 3.1

Characteristics of store locations

BELGIUM	Average	Std. Dev.	Minimum	Maximum
GDP per capita (€)	35.106,58	10.524	15.700	63.330
Population Density	1.190,93	2.302,45	36,27	16.393,32
N° of Competitors	51,48	43,26	3	225
Next to the Border (20km)	45,40%	49,86	0	1
FRANCE	**Average**	**Std. Dev.**	**Minimum**	**Maximum**
GDP per capita	28.828,17	5.796	20.400	42.500
Population Density	378,18	650,88	9,25	4.635,45

Table 3.1 above provides some descriptive statistics about the store locations. We use a set of proxies to control for some supply-side and demand-side factors that could explain spatial heterogeneity in the tax pass-through. To measure the intensity of competition faced by each store, we use a variable indicating the number of competing retailers within a driving distance of 15 minutes. These data are collected by a private company that provides contact information to suppliers about supermarkets and grocery stores located in Belgium. From their postal address, it is then possible to compute the driving distance from each store to any other retailer in the area. However, this variable is only available for Belgian stores. Therefore, we cannot

[46] We also estimate the models using the average monthly price to check whether including temporary price discounts affects our results. Yet, this exercise still confirms our findings. These results are available upon request.

directly control for competitive pressure in French stores. To check for the robustness of our results, we will use local density of population (in quartile) as a proxy for competitive pressure. Thus, we compare the evolution of prices between Belgian and French stores that are in the same quartile of the population density distribution of their respective country. Using each store geographical location, we can also compute their distance to the nearest border. This enables checking whether those stores close to the border (subject to potential cross-border shopping) responded differently to the tax change. Furthermore, to control for demand-side local heterogeneity, each store is matched with the average GDP per capita at the Local Administrative Unit Level (NUTS 3) and population density data at the municipality level.

3.5 The Empirical Models

In order to estimate the tax pass-through to spirits' retail prices, we perform a Difference-in-Differences analysis separately on six distinct products, by considering the retail prices of the same products sold in France as a control group. The use of French prices for the same brand as a counterfactual can potentially control for unobserved factors, common to both France and Belgium, that could have affected the brand retail price over the period of policy implementation. The analysis is organized as follows. Firstly, we estimate for each brand the tax pass-through at the chain level. This gives us a measure of how the tax was shifted across retail stores on average. Secondly, we estimate for each brand the tax pass-through at the store level. This exercise allows assessing the degree of tax pass-through heterogeneity across different geographical locations. We test whether such heterogeneity is associated to differences in local competition and/or proximity to the border. Lastly, we account for time heterogeneity in order to see how the tax shifting evolved during the period. These estimates are also important to check whether the spatial variation in tax pass-through was permanent or just temporary.

All models are estimated using the standard OLS procedure. A main concern in the difference-in-difference literature is that errors can be correlated across different groups of observations. In that case, assuming that errors are independent across observations can lead to an incorrect estimation of the standard errors for the treatment effects (Bertrand, Duflo & Mullainathan, 2003). In our context, the potential sources of correlation are: (i) serial correlation of errors for each store; and (ii) spatial correlation of errors across stores. The first one is standard when observing the same individual/firm over multiple periods and it can be produced by unobserved characteristics that are constant overtime. The second one can be produced by local shocks that affect st=ores in the same area similarly. This source of correlation is quite relevant in our case since stores set their prices by matching the lowest price of any competitors within a certain radius. To account for these two possible sources of error correlation, we cluster errors at the arrondissement level. As a result, we use around 60 clusters for each product.[47] This allows us to account for both serial correlation of errors for each store and shocks that could affect stores in the same area equally. Each model is estimated separately for each of the six products analyzed.

3.5.1 Average Tax Pass-Through

In this section, we estimate the average tax pass-through to the retail price of each spirit considered. We use the standard difference-in-differences procedure. The retail price for each specific brand in store i during month t is expressed as follows.[48]

$$P_{it} = \beta_0 + \beta_1 BE_i + \beta_2 T_t + \beta_3 (BE_i \times T_t) + \varepsilon_{it}. \tag{3.1}$$

β_0 is the pre-reform price level in France. While BE_i is a dummy variable equal to 1 if the store i is located in Belgium and 0 if located in France. Its coefficient β_1 measures the pre-reform difference in prices between Belgium and France. The variable T_t is a dummy variable equal to 1 during the period of tax implementation (post November 2015) and 0 otherwise. Its coefficient β_2 measures the price difference between the pre-reform and post-reform period in France, which serves as a counterfactual for the price evolution in Belgium. The fourth term is the interaction of the treated group BE_i and the post-reform variable T_t. Its coefficient β_3 captures the price increase in Belgium due to tax change and allows computing the tax pass-through rate as follows:

$$Tax\ Pass\ Through\ Rate = \frac{\beta_3}{\Delta tax} \times 100.$$

This work focusses on the short-run impact of the tax on retail prices, with a narrow time window going from August 2015 until March 2016. In this a way, we actually compute the difference in the average price of the product in Belgium between the three months period before the tax reform (August 2015 - October 2015) and the five months period after the tax reform (November 2015 - March 2016). This price change in the treated group (stores in Belgium) is then compared with the price change of the same product between the two periods in the control group (stores in France). A fundamental assumption, however, is that nothing else a part from the tax should have affected the retail price for the same spirits' brand in Belgium

[47] We also run the models clustering at either store, province, or country level. In every case, we find smaller standard errors. Thus, we are reporting the most conservative estimates (i.e. those with the largest standard errors).

[48] The brand index is dropped in the rest of the analysis to ease notation.

and France differently in the period after the tax implementation. As the period is quite narrow, it is quite easy to check that there was no major policy change in Belgium and France that should have impacted the product prices in the two countries.

Table 3.2

Average Tax Pass-through (Model 3.1)

	Product					
	A	**B**	**C**	**D**	**E**	**F**
Intercept (β_0)	16,31*** (0,08)	11,77*** (0,04)	12,88*** (0,08)	14,36*** (0,06)	15,03*** (0,09)	14,80*** (0,11)
Treated (β_1)	-0,71*** (0,08)	-0,50*** (0,05)	-2,10*** (0,09)	-0,84*** (0,06)	-0,85*** (0,09)	0,22** (0,11)
Post-reform (β_2)	-0,10** (0,05)	-0,09*** (0,03)	-0,14 (0,10)	0,10 (0,07)	-0,06 (0,07)	-0,17** (0,07)
Treatment (β_3)	**3,30***** (0,05)	**2,67***** (0,05)	**2,80***** (0,11)	**2,64***** (0,08)	**2,54***** (0,09)	**3,13***** (0,07)
N° Observations	2960	3096	3248	3256	3240	3208
Product type	Vodka	Vodka	Vodka	Whiskey	Rum	Rum
% Alcohol	40%	37,5%	37,5%	40%	37,5%	40%
Excise Tax increase	2,43€	2,28€	2,28€	2,43€	2,28€	2,43€
% Pass-Through	**135,80**	**117,11**	**122,81**	**108,64**	**111,40**	**128,81**
Confidence Interval	131,68 - 139,91	112,28 - 121,49	113,60 - 132,02	102,06 - 115,64	103,51 - 119,74	122,63 – 134,98

Notes: *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the arrondissement level, are in parenthesis.

Table 3.2 above shows the estimated coefficients of *model 3.1*. The first line of Table 3.2 shows the intercept of the model for each product, which indicates the average product price in France in the pre-tax period. The line "Treated" shows how prices in Belgium (treated group) differ from France (control group) before the reform. The "Post-reform" line displays the price evolution in France after the reform (November 2015). Most of these coefficients are slightly negative and close to zero, thus suggesting as counterfactual that spirits prices would have slightly declined in Belgium without the tax increase. Yet, just three of them are statistically significant at the 5% level. The line "Treatment" shows the impact of the tax reform on the Belgian price for each product. These coefficients can be interpreted as the price change in € induced by the tax reform. As the products considered differ in their alcohol content, the tax change was different across products. From the tax hike specific to each product and its treatment coefficient β_3, it is then possible to calculate the tax pass-through rate. As shown in Table 3.2, the tax pass-through rate tends to vary across products. The tax was over-shifted to the retail prices of all spirits with a confidence level of 95%. This cross-product variation in pass-through can be related to supply side and demand side differences across products. We will not explore further this cross-product variation in the tax pass-through. Instead, we will study the spatial variation in the tax pass-through for each product separately.

3.5.2 Spatial Heterogeneity in the Tax Pass-Through

In this section, we focus on identifying the spatial variation of tax pass-through for the same product across stores. To get a preliminary measure of heterogeneity in tax shifting, we compare spatial price dispersion in both Belgium and France before and after the tax reform. The spatial price variance of each spirit across Belgian stores has significantly increased after the tax reform, while it stayed constant over the same period in France. A Levene's Test on the homogeneity of spatial price variances between the pre-reform and post-reform period is rejected for all products in the treated group with the 99% confidence level (except for F). While it is accepted for all products in the control group (except F, for which it has slightly declined).[49]

To provide more compelling evidence about the evolution of spatial differences in spirit prices, we estimate the same model as above (model 3.1) by including both store fixed effects and a store specific treatment effect. This will deliver a store specific tax pass-through. Store fixed effects are fundamental in order to capture tax pass-through heterogeneity. This is because they can account for possible pre-reform (time invariant) unobserved factors that affect the store's pricing. These can include differences in the cost of selling the products (such as transportation costs, rents or local wages) and in price elasticity of demand. If we do not correctly control for these pre-reform differences in prices across stores, there is a risk of confounding them with heterogeneity in tax-shifting. From now on, every model we present includes store fixed effects. Formally, we estimate the following regression for each product:

$$P_{it} = \delta_i + \beta_2 T_t + \beta_{3i}(BE_i \times T_t \times \delta_i) + \varepsilon_{it}. \tag{3.2}$$

Where δ_i are the fixed effects coefficients for each store i located in either Belgium or France. These are captured by store-specific dummy variables and give the average price level of each store i before the tax reform. The coefficient β_2 is capturing the evolution of the average price in French stores after the tax reform. Which is seen as the counter-factual scenario. While β_{3i}

[49] The results of this test can be found in Table 3.A.1 in the appendix.

is the store i's specific tax pass-through if this store is located in Belgium. The results of these estimations are shown in the figures below (from 3.3.A to 3.3.F). Results are aggregated at the municipality level. Every color represents a certain degree of tax pass-through in a given municipality. Interestingly, since these stores are local price followers, their tax shifting should be indicative of the general trend in spirit price changes for each geographical location. These figures display heterogeneous tax shifting across space after the tax reform. Although the tax was over-shifted to different extents in most municipalities, there are also some areas where the tax was instead under-shifted (blue areas in the figures).

Figure 3.3.A: Product A

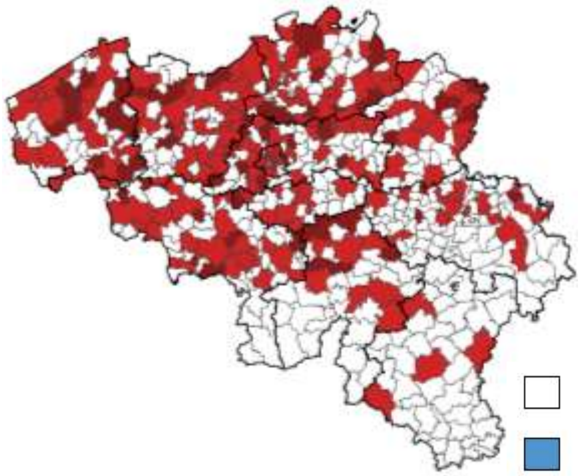

Figure 3.3.B: Product B

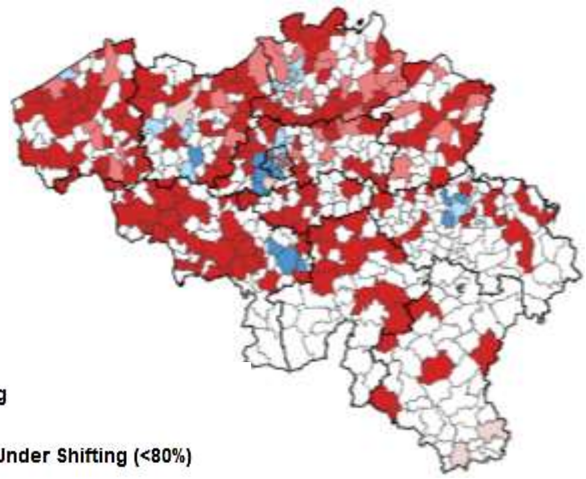

Figure 3.3.C: Product C

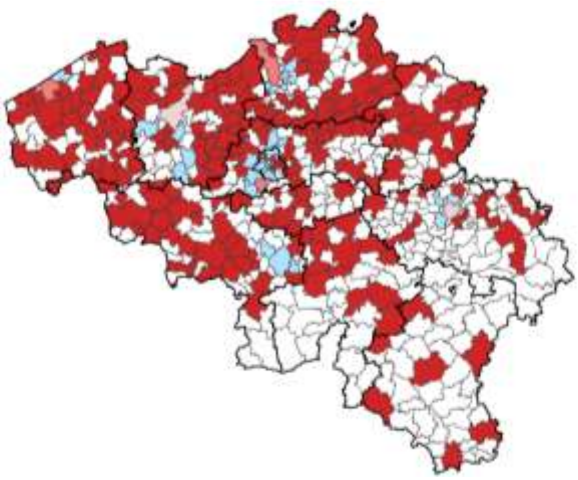

Figure 3.3.D: Product D

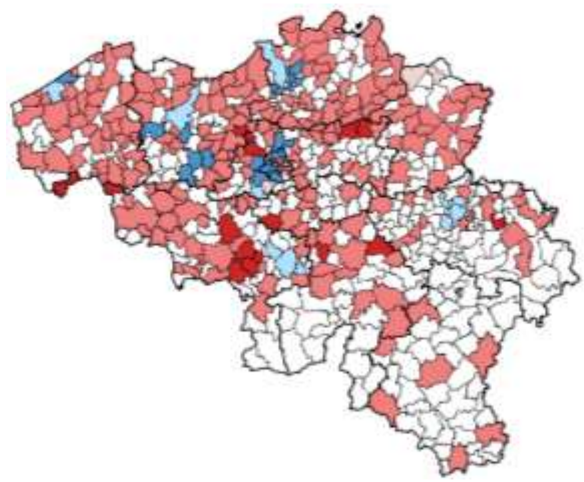

Figure 3.3.E: Product E

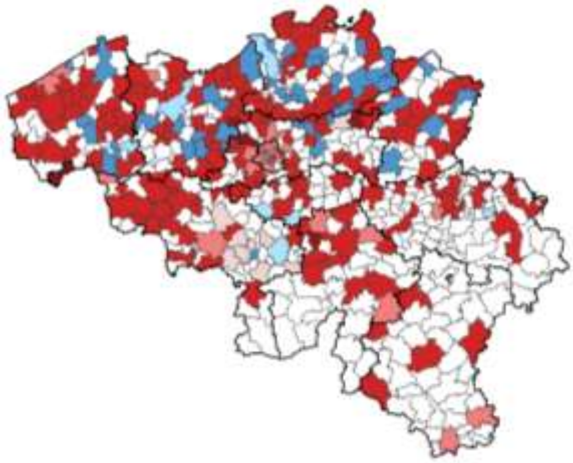

Figure 3.3.F: Product

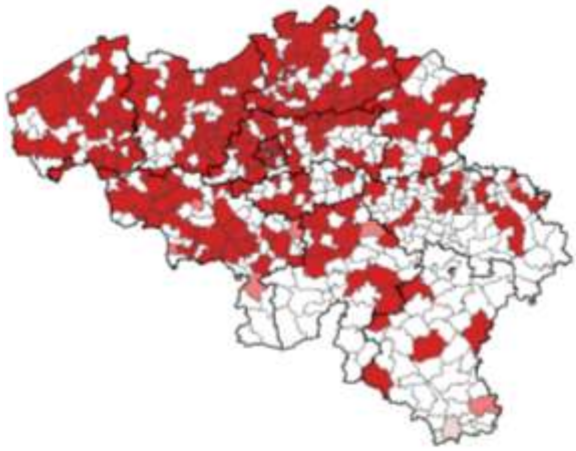

Variation in the tax pass-through is related to variation in market structure and price elasticity of demand. Thus, accounting for spatial differences in these two factors can enable us to understand such heterogeneity in tax shifting. In order to do so, we proceed as follows. First, we test for the effect of local competition on the tax pass-through at the store level. To account for local differences in market structure, the model contains information about the intensity of competition at the store level. Intuitively, one would expect lower tax pass-through when there are many competitors nearby. Second, we focus on the proximity to the border. The scope for cross-border shopping may be quite important in Belgium, a relatively small country, because a large part of the population lives in proximity to the border (and there are many cross-border workers). This is also relevant because Belgium shares borders with several different countries which set different alcohol taxes. For this reason, we also estimate a model that includes the proximity to the border of each store. That model allows us to test for differences in price setting for stores close to the border. If cross-border shopping is an effective threat for those stores, tax shifting in border areas should be lower as the demand elasticity would be higher. Third, as demand side factors may distort our results, we also estimate a model that includes information about spatial heterogeneity in some supply-side and demand-side factors.

3.5.3 Intensity of Competition

Having information about the number of competing retailers for each store allows us to test for the effect of competition on the tax pass-through. As we are comparing the tax shifting of the same product across different geographical locations, it is clear that we restrict our focus to the intensity of competition among retailers and not among producers. Each product analyzed is among the world's most popular brands in their respective category and none of their producers is vertically integrated with any Belgian or French retailer. To test whether the local intensity of competition at the store level can be related to the observed spatial heterogeneity in tax pass-through, we compare the tax shifting among areas exhibiting a low, medium or high intensity of competition. We define the intensity of competition in terms of number of local competitors for each store within a driving distance of 15 minutes. The competitors are from different supermarket chains than the chain under study. A store is considered in a low-competition cluster if it falls in the first quartile of this distribution with at most 26 local competitors. A store is in a medium-competition cluster if it falls in the 2nd or 3rd quartile of the distribution with between 27 and 59 local competitors. While it is in a high-competition cluster if it is in the last quartile of the distribution with more than 60 competitors. Formally, we estimate the following regression:

$$P_{it} = \delta_i + \beta_2 T_t + \beta_L \left(BE_i \times T_t \times Low_{Comp_i}\right) +$$

$$+\beta_M \left(BE_i \times T_t \times Med_{Comp_i}\right) + \beta_H \left(BE_i \times T_t \times High_{Comp_i}\right) + \varepsilon_{it}. \quad (3.3)$$

Where Low_{Comp_i}, Med_{Comp_i} and $High_{Comp_i}$ are dummy variables equal to one if store i is in either a low, medium or high competition cluster. We want to estimate the coefficients β_L, β_M and β_H, reflecting the tax-pass through specific to each of these three competition clusters. We expect these coefficients to be statistically different from each other and, in particular, to decrease with the intensity of competition. That is, we expect to find that $\beta_L > \beta_M > \beta_H$. The results of this estimation are displayed in Table 3.3 below. The last two rows of this table also show the results of the Wald test on the equality of coefficients for low and high competition. Where the null hypothesis is that there is no difference in tax shifting between low and high competition. That is, H_0: $\beta_L = \beta_H$.

The results of Table 3.3 tend to confirm our theoretical prediction. The price increase was smaller in high competition areas. The magnitude of this effect, however, can vary across products. For most products, the difference in tax shifting between low and high competition is between 0,40€ and 0,50€. The magnitude of such effect is much smaller for product A, for which this difference is equal to 0,11€. While it is absent for product F. The test on the equality of coefficients for high and low competition indicates that, except for product F, these differences in tax shifting are statistically significant at the 99% confidence level. Therefore, the results of *model 3.3* suggest that the tax shifting decreased with the intensity of competition at the local level.

Table 3.3

Tax Pass-Through and Intensity of Competition (Model 3.3)

	Product					
	A	**B**	**C**	**D**	**E**	**F**
Low Competition	**3,36**	**2,82**	**2,92**	**2,76**	**2,79**	**3,11**
$(\boldsymbol{\beta_L})$	(0,06)	(0,04)	(0,11)	(0,08)	(0,09)	(0,08)
Medium Competition	**3,29**	**2,78**	**2,91**	**2,78**	**2,48**	**3,15**
$(\boldsymbol{\beta_M})$	(0,05)	(0,04)	(0,11)	(0,08)	(0,12)	(0,08)
High Competition	**3,25**	**2,32**	**2,47**	**2,27**	**2,41**	**3,11**
$(\boldsymbol{\beta_H})$	(0,06)	(0,08)	(0,08)	(0,12)	(0,14)	(0,08)
	Test on the Equality of Coefficients $(\boldsymbol{H_0: \beta_L = \beta_H})$					
F value	13,98	46,78	39,05	26,12	8,76	0,04
p-value	<0,01	<0,01	<0,01	<0,01	<0,01	0,84

Notes: All coefficients are statistically significant at the 0,01 level. Standard errors, clustered at the arrondissement level, are in parenthesis. The last two rows show the results of the Wald test on the equality of the coefficients for low and high competition, where the null hypothesis is $H_0: \beta_L = \beta_H$.

To retrieve the tax pass-through rate for each competition level, we divide the treatment coefficients presented in Table 3.3 by the product specific increase in the excise tax. The results are displayed above in Table 3.4. As already suggested in Table 3.3, the tax pass-through rate varies with the intensity of competition. The tax was largely over-shifted with low competition. Whereas, it was shifted to a lesser extent or even under-shifted with high competition. This indicates that the extent of tax shifting, and the intensity of competition are indeed negatively correlated.

Table 3.4

Tax Pass-Through Rate and Intensity of Competition

	Product					
	A	**B**	**C**	**D**	**E**	**F**
Low Competition	**138%**	**124%**	**128%**	**114%**	**122%**	**128%**
(C.I.)	133-142	120-127	119-137	107-121	117-131	121-135
Medium Competition	**135%**	**122%**	**128%**	**114%**	**109%**	**130%**
(C.I.)	130 -140	118-125	118-138	108-121	98-119	123-136
High Competition	**134%**	**102%**	**108%**	**93%**	**105%**	**128%**
(C.I.)	129-138	94-108	97-119	84-103	94-117	121-135

3.5.4 Cross-Border Shopping

Another possible source of the tax pass-through heterogeneity is the proximity to the border. Cross-border shopping can be quite important in Belgium since a large part of the population lives close to the border. In our sample, 45,4% of Belgian stores are within a distance of 20km to the border. Moreover, Belgium shares borders with four different countries (France, Luxembourg, Germany and The Netherlands), which have different levels of alcohol taxation and spirit prices. The alcohol tax reform in Belgium has considerably increased the price gap in spirit prices between Belgian and foreign stores. Luxembourg and to a lesser extent Germany, had lower spirit prices before the reform. Whereas the Netherlands and to a lesser extent France, had higher spirit prices before the reform. In order to investigate the relationship between tax pass-through and the scope for cross-border shopping, we estimate a model that includes information about the proximity to the border of each store. This allows testing for differences in tax shifting according to whether or not stores are close to the border. For each specific product, we estimate the following model:

$$P_{it} = \delta_i + \beta_2 T_t + \beta_3 (BE_i \times T_t) + \beta_{BR}\left(BE_i \times T_t \times BR_{km_i}\right) + \varepsilon_{it}. \tag{3.4}$$

The only difference here is the inclusion of the last interaction term: $\left(BE_i \times T_t \times BR_{km_i}\right)$. Where BR_{km_i} is a dummy variable indicating whether store i is within a certain km distance to the border. The coefficient β_{BR} therefore measures the difference in the treatment effect (tax shifting) for those stores that are within that certain distance to the border. In particular, we use three different distances. Namely 10km, 15km or 20km. As long as cross-border shopping is really binding price decisions, we expect β_{BR} to be negative and significantly different from zero.

The results of *model 3.4* are displayed in table 3.5 below. Table 3.5 shows that tax shifting did not change with the proximity to any border.[50] At any distance considered, those stores close to the border did not shift differently the tax to the retail price compared to other stores. We obtain the same results even when controlling for the intensity of competition as in *model 3.3*. This suggests that the threat of cross-border shopping does not seem to play a significant role in the shifting of the tax on spirit prices, even though the price gap with several neighboring countries increased substantially after the reform. A possible explanation for this can be the fact that the

[50] Although we find a slightly positive difference for stores within 10km distance from the border for two products, this disappears once controlling for the number of competitors.

price gap with neighboring countries was not high enough to justify a price adjustment at the border or that Belgian stores are poorly informed about foreign prices near the border. Another possible option could be the market segmentation between mobile and immobile shoppers. The stores locate close to the border only retain the non cross-border shoppers (immobile shoppers) who are likely to exhibit more inelastic demand than the cross-border shoppers (mobile shoppers). This effect could offset the downward pressing effect of cross-border shopping on prices.

Table 3.5

Tax Pass-Through and Proximity to any Border (Model 3.4)

	Product					
$\boldsymbol{\beta_{BR}}$	**A**	**B**	**C**	**D**	**E**	**F**
Border at 20 Km	-0,01 (0,03)	0,07 (0,11)	0,09 (0,09)	0,07 (0,09)	0,03 (0,11)	-0,01 (0,01)
Border at 15 Km	0,03 (0,03)	0,07 (0,09)	0,12 (0,08)	0,12 (0,09)	0,14 (0,11)	-0,03 (0,01)
Border at 10 Km	0,06** (0,03)	0,06 (0,09)	0,11 (0,09)	0,13 (0,10)	0,22** (0,10)	-0,02 (0,02)

Notes: *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the arrondissement level, are in parenthesis. Each row shows the estimated coefficient β_{BR} for every product considering stores within a 10km, 15km or 20km distance to any border.

The absence of border effect on tax shifting may also be due to the averaging out of various border effects among the four different neighboring countries. Indeed, if the border effect depends on the size and the sign of the price gap, we may expect different border effects for the four different countries, notably for Luxembourg with the lowest spirit price. We test for this hypothesis by re-estimating *model 3.4* differently. That is, we now consider each border separately to estimate how tax shifting varies when a store is close to a specific border. In doing so, we did not find any significant impact when considering just those stores at the border with either France, Netherlands or Germany. Where prices were respectively comparable, higher or slightly lower than in Belgium before the reform.[51] However, we did find some interesting results for those stores close to Luxembourg (where spirits were on average 4€ cheaper before the tax reform).

[51] As for *model 4.1*, no effect was found when considering stores within either 20km, 15km or 10km from the border.

In our sample, we only have three stores that are located within 10km distance from the Luxembourg border and no other store is located within 20km. These stores are all located in remote areas with a small number of competitors (less than nine) and hence they face a quite low competition. As we have learned from the results of *model 3.3*, this means that the tax shifting of these stores should be significantly higher than the one of stores facing more competition. Yet, if competition at the Luxembourg border matters, this effect can be ambiguous. This is because the lower domestic competition could be offset by the higher foreign competition from Luxembourg. In order to limit cross-border shopping, these stores could have shifted the tax on spirit prices to a lesser extent compared to those stores facing a similar domestic competition but no proximity to the border. Formally, to measure the tax pass-through of stores at the border of Luxembourg we estimate the following regression for each product separately:

$$P_{it} = \delta_i + \beta_2 T_t + \beta_L \left(BE_i \times T_t \times Low_{Comp_i} \times NoLUX_{B_i}\right) + \\ +\beta_{LUX}\left(BE_i \times T_t \times Low_{Comp_i} \times LUX_{B_i}\right) + \beta_M \left(BE_i \times T_t \times Med_{Comp_i}\right) + \\ +\beta_H \left(BE_i \times T_t \times High_{Comp_i}\right) + \varepsilon_{it}. \tag{3.5}$$

Where Low_{Comp_i}, Med_{Comp_i} and $High_{Comp_i}$ are the same variables as in *model 3.3*. However, the first interaction term includes the dummy variable $NoLUX_{B_i}$, which is equal to 1 if store i is not at the border of Luxembourg (within 10km). The coefficient β_L therefore measures the tax pass-through of stores facing low competition and not at the border of Luxembourg. The dummy variable LUX_{B_i} is instead equal to 1 if a store is close to Luxembourg (within 10km). Hence, the coefficient β_{LUX} measures the tax pass-through of these stores, which are also all facing low domestic competition. The other variables are the same as in *model 3.3*. The objective of this regression is to estimate β_{LUX} and test whether $\beta_{LUX} < \beta_L$. That is, we would like to know whether for the same level of (domestic) competition, tax shifting decreases with the proximity to the border of Luxembourg.

The results of *model 3.5* are displayed in Table 3.6 below. From this table we can compare the tax pass-through of store close to Luxembourg (β_{LUX}) with other stores located in low competition areas (β_L) Interestingly, the tax pass-through of stores close to Luxembourg seems to be lower than the one of other stores in low competition areas. This is true for most product. Yet, the Wald test on the equality of coefficient suggests that only three of these differences in tax pass-through are significant at the 0,05 level. These are the products A, B and F. For product B and F, such difference is quite large, being close to 0,40€, while it is small for product A, being only 0,06€. The difference is 0,17€ for product E, but it is only significant at the 0,10 level. This heterogeneity in the "border effect" across products might depend on many factors, such as different tastes for different products to make it worth doing cross border shopping or the effective supply of those same products on the other side of the border. This cross-product heterogeneity of the "border effect" also suggests that it is important to analyze the tax pass-through at the product level. Since we could not have found this effect when averaging out the border effect over different products.

Table 3.6

Tax Pass-Through and Proximity to Luxembourg (Model 3.5)

	Product					
	A	**B**	**C**	**D**	**E**	**F**
Low Comp. and No Prox. to Lux. ($\boldsymbol{\beta_L}$)	**3,36** (0,06)	**2,83** (0,04)	**2,92** (0,11)	**2,76** (0,08)	**2,80** (0,10)	**3,12** (0,08)
Low Comp. and Prox. Lux. ($\boldsymbol{\beta_{LUX}}$)	**3,30** (0,05)	**2,45** (0,18)	**3,00** (0,11)	**2,70** (0,09)	**2,63** (0,11)	**2,73** (0,19)
Medium Competition ($\boldsymbol{\beta_M}$)	3,29 (0,06)	2,78 (0,04)	2,91 (0,11)	2,78 (0,08)	2,48 (0,12)	3,15 (0,08)
High Competition ($\boldsymbol{\beta_H}$)	3,25 (0,06)	2,32 (0,08)	2,47 (0,13)	2,27 (0,12)	2,41 (0,14)	3,11 (0,08)
	Test on the Equality of Coefficients ($\boldsymbol{H_0: \beta_L = \beta_{LUX}}$)					
F value	15,49	4,42	3,10	2,09	3,10	4,90
p-value	<0,01	0,04	0,08	0,15	0,08	0,03

Notes: All coefficients are statistically significant at the 0,01 level. Standard errors, clustered at the arrondissement level, are in parenthesis. The last two rows show the results of the Wald test on the equality of the coefficients for low competition areas either close to (β_{LUX}) or far away (β_L) from Luxembourg, where the null hypothesis is H_0: $\beta_L = \beta_{LUX}$.

The results of *model 3.4* and *model 3.5* suggest that only a significant price gap with a neighboring country can reduce tax shifting for some products (but not for all). This is confirming the standard view that the scope for cross-border shopping increases with the price gap between neighboring countries. Yet, the absence of "border effect" for stores close to either France (where spirit prices were only 0,5€ higher before the tax) or Germany (where spirit prices were around 1€ lower before the tax) could also suggest a lack of information/attention about foreign prices.

3.5.5 Demand-side Heterogeneity

All models estimated so far provide a supply-side explanation on the spatial heterogeneity in the tax pass-through based on the idea that domestic and foreign competition circumstances vary across space. Yet, tax incidence can also depend on the demand circumstances that may also vary across space. Therefore, we estimate another model of tax pass-through heterogeneity that controls for some differences in demand-side characteristics. We do that by including information about local population density (whether the store is in a rural area or not) and the GDP per capita at the arrondissement level. Intensity of competition is now measured by the log of competing stores within a driving distance of 15 minutes from the store. We account for proximity to the border as in *model 3.4.* The treatment coefficient estimates of this model will tell us whether the heterogeneity in the tax pass-through is still correlated to the intensity of competition and proximity to Luxembourg after controlling for differences in some observable demand-side characteristics (such as rural/urban status and GDP per capita). For each specific product, we estimate the following regression:

$$
\begin{aligned}
P_{it} = \delta_i + \beta_2 T_t + \beta_3 (BE_i \times T_t) + \beta_{Y_F}(T_t \times \ln(Y)_i) + \beta_{Y_B}(BE_i \times T_t \times \ln(Y)_i) + \\
+\beta_{R_F}(T_t \times Rural_i) + \beta_{R_B}(BE_i \times T_t \times Rural_i) + \\
+\beta_C(BE_i \times T_t \times \ln(COMP)_i) + \beta_{LUX}\left(BE_i \times T_t \times LUX_{B_i}\right) + \varepsilon_{it}. \qquad (3.6)
\end{aligned}
$$

As for every other specification, δ_i is the store specific fixed effect, which captures all those pre-reform unobserved factors that are store specific and time-invariant. The coefficients β_2 and β_3 measure the baseline of the counterfactual and the treatment effect respectively. The variable $\log(Y)_i$ is the log of the GDP per capita in the arrondissement in which store i is located. While $Rural_i$ is a dummy variable equal to one when the store is in a rural area (with less than 200 inhabitants per km^2). Each of these variables is interacted with the post reform dummy (T_t) and the treatment interaction term ($BE_i \times T_t$). Their respective coefficients measure how prices evolved after the reform in the control (France) and in the treated group (Belgium). In particular, β_{Y_B} and β_{R_B} measure the additional effect of treated stores relative to control stores in areas with higher GDP and rural areas, respectively. $\ln(COMP)_i$ is the log of the number of competing retailers for store i. The coefficient β_C measures how tax shifting varies with the number of competing retailers. If results of *model 3.3* are confirmed, we expect to find $\beta_C < 0$. That is, tax pass-through should decrease with competition. LUX_{B_i} is a dummy variable indicating if a store is close to the border with Luxembourg (within 10km). Because in this model we include the baseline treatment effect, the interpretation of β_{LUX} is slightly different from the one of *model 3.5*. Here β_{LUX} estimates directly by how much tax shifting differs in these areas with respect to the average store, once controlling for some spatial differences in demand-side characteristics and the number of competing retailers.

Table 3.7

Controlling for Demand-side Characteristics (Model 3.6)

	Product					
	A	**B**	**C**	**D**	**E**	**F**
"Gross" Treatment ($\boldsymbol{\beta_3}$)	-2,83 (1,72)	3,65 (2,23)	2,12 (5,04)	-2,49 (3,46)	-7,05** (3,50)	-3,71* (3,16)
GDP per capita FR ($\boldsymbol{\beta_{Y_F}}$)	-0,54*** (0,17)	0,04 (0,19)	-0,10 (0,55)	-0,54* (0,33)	-0,85*** (0,28)	-0,65** (0,31)
GDP per capita BE ($\boldsymbol{\beta_{Y_B}}$)	0,62*** (0,17)	0,01 (0,23)	0,18 (0,51)	0,62* (0,34)	0,98*** (0,35)	0,68** (0,31)
Rural areas FR ($\boldsymbol{\beta_{R_F}}$)	0,09 (0,10)	0,04 (0,05)	0,26 (0,16)	0,24* (0,12)	0,00 (0,08)	0,05 (0,04)
Rural areas BE ($\boldsymbol{\beta_{R_B}}$)	0,00 (0,11)	-0,08 (0,07)	-0,33* (0,16)	-0,32** (0,14)	0,27** (0,11)	-0,10 (0,05)
N° of Competitors ($\boldsymbol{\beta_C}$)	**-0,06***** (0,02)	**-0,30***** (0,06)	**-0,27***** (0,07)	**-0,30***** (0,08)	**-0,15*** (0,08)	-0,02 (0,01)
Proximity to Luxembourg ($\boldsymbol{\beta_{LUX}}$)	**-0,16***** (0,04)	**-0,72***** (0,18)	**-0,21**** (0,09)	**-0,43***** (0,12)	**-0,27*** (0,14)	**-0,40**** (0,18)

Notes: *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the arrondissement level, are in parenthesis.

The estimates of *model 3.6* are reported in Table 3.7 above. The coefficients β_3 is the "gross" treatment effect. Although it is negative for most products, that does not mean a net negative treatment effect. Indeed, one must take into account the other treatment interaction effects, notably the coefficient β_{Y_B} for the GDP interaction that is positive for every product, although not always significant. Consider for instance product E. Its "gross" treatment effect β_3 is equal to -7,05, while β_{Y_B} amounts to 0,98. Considering that the lowest GDP per capita amounts to 15.700€, taking the log and multiplying by the β_{Y_B} we obtain $\ln(15.700) \times 0{,}98 = 9{,}47$. The net treatment effect after controlling for the GDP is then equal to 9,47 - 7,05=2,42. This indicates that in areas with the lowest GDP per capita, prices in Belgium after the tax reform increased by 2,42€ more than similar areas (in terms of GDP) in France. As all other stores have a higher GDP per capita, the treatment effect after controlling for GDP must be greater than this figure.[52] The fact that β_{Y_B} tend to be positive for most products suggests that spirit prices in Belgium increased by more in richer areas compared to France. Yet, the results in Table 3.7 also show that this effect is mostly driven by a decline of spirit prices for stores located in richer areas in France. Furthermore, stores in rural areas do not seem to follow any particular trend after the reform.

Interestingly, the results of *model 3.6* seem to confirm our previous findings on the correlation between the tax pass-through and the local competition. The extent of the tax pass-through is negatively correlated to the number of local competitors for all products except F. This effect is more prevalent and it is similar in magnitude for products B, C and D. It is smaller but still significantly different from zero for product A, while it is only significant at the 10% level for product E. To get an idea on the magnitude of the competition effect on tax shifting, we compute how the tax pass-through changes when increasing the number of competitors from 20 to 100. We consider the case of a store located in an area with the average GDP per capita and ignore the rural area and border effect. Considering product D, the treatment effect for a store with only 20 competitors would be equal to:

$$\tau_{20} = \beta_3 + \left(\ln(Y_i) \times B_{Y_B}\right) + (\ln(COMP_i) \times B_C) = -2{,}49 + (\ln(35.100) \times 0{,}62) - (\ln(20) \times 0{,}30) = 3{,}03.$$

While if the number of competitors rises to 100 we get:

$$\tau_{100} = -2{,}49 + (\ln(35.100) \times 0{,}62) - (\ln(100) \times 0{,}30) = 2{,}55.$$

Which means that increasing the number of competitors from 20 to 100 decreases the tax shifting by 0,48€. These results are in line with those of *model 3.3*, in which the difference in tax shifting between low and high competition areas for product D was 0,49€ on average.

Model 3.6 also confirms that stores close to Luxembourg tend to set lower spirit prices after the tax reform as β_{LUX} is negative and significant for most products. This effect seems more pronounced than the one found in *model 3.5.* Although the two coefficients have a different interpretation and cannot be directly compared. This is probably because *model 3.6* controls for the number of competing retailers through a continuous variable (i.e. the natural log of the number of competitors), which is extremely low at the Luxembourg border (less than nine). Hence, the Luxembourg border dummy could also capture some non-linearity in the relationship between the number of competitors and the tax shifting. Overall, these results indicate that, after controlling for some observable heterogeneity in demand-side

[52] However, in order to compute the overall net treatment effect all other treatment interaction terms must also be taken into account.

characteristics, the number of competing retailers and proximity to Luxembourg (the lowest price country) are still significantly correlated with the heterogeneity in the tax shifting.

3.5.6 Robustness Checks

A possible concern in estimating the relationship between competition and tax pass-through can be the lack of a proper counterfactual for stores facing a similar degree of competition in France (our control group). As we do not observe the number of competitors for the French stores, we did not formally check whether spirit prices in France have changed differently after the reform in high competition and low competition areas. The validity of the control group requires to compare stores in France and in Belgium facing the same level of competition. The results of the Levene's test presented at the beginning of this section show that the spatial price dispersion was mostly stable in France after the tax reform, while it increased substantially in Belgium. This suggests that the spirit prices in the control group did not diverge much across stores facing different competition after the tax reform. However, it is still possible that this "average effect" conceals heterogeneous changes between high competition and low competition stores in France.

To address this issue, we run another model using population density at the local level (municipality) as a substitute to proxy for the intensity of competition. In such a way, we can compare stores facing different intensity of competition (proxied by the population density) both in France (control group) and in Belgium (treated group). The assumption here is that French stores face more competition in high population density areas. The use of population density to measure the intensity of competition at the local level is not a bad proxy. As shown in Figure 3.4 below, the number of stores in a local area is highly correlated with the population density in Belgium.

Figure 3.4

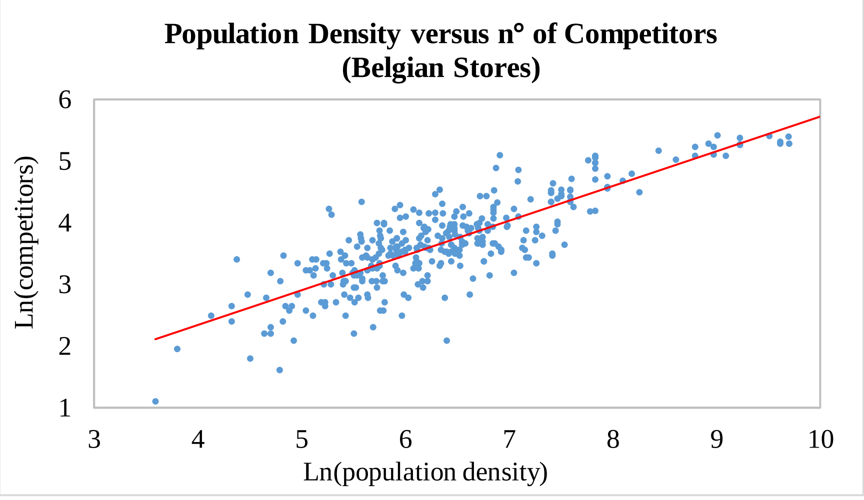

The idea is to re-estimate *model 3.5* by using the population density at the municipality level instead of the number of competitors. To control for the difference in population density among Belgian and French municipalities we will express the population density in quartiles in the regression. In such a way, we compare the price changes between Belgian and French stores that are in the same quartile of the population density distribution of their respective country. For instance, we consider in the low competition areas, those stores that are in the first quartile of the population density distribution of either Belgium or France. Formally, for each product we estimate the following model:

$$P_{it} = \delta_i + \beta_{L_F}(T_t \times Low_{den_i}) + \beta_{L_B}(BE_i \times T_t \times Low_{den_i} \times NoLUX_{B_i}) +$$

$$+\beta_{LUX}(BE_i \times T_t \times Low_{den_i} \times LUX_{B_i}) + \beta_{M_F}(T_t \times Med_{den_i}) + \beta_{M_B}(BE_i \times T_t \times Med_{den_i}) +$$

$$+\beta_{H_F}(T_t \times High_{den_i}) + \beta_{H_B}(BE_i \times T_t \times High_{den_i}) + \varepsilon_{it}. \tag{3.7}$$

The structure of *model 3.7* is similar to *model 3.5.* Here the difference is that we use population density as a proxy for competition so that we can control for the price changes of French stores facing different level of competition. The counterfactual scenarios for different levels of competition are captured by the coefficients β_{L_F}, β_{M_F} and β_{H_F}. Which correspond to the post-reform price changes in France for stores that are in low, medium or high competition areas, respectively. The coefficients β_{L_B} and β_{LUX} measure the tax pass-through for Belgian stores in low competition areas not close and close to Luxembourg, respectively. Note that their counterfactual scenario is not the same as in *model 3.5*, where we use the average price change in France β_2. Here the counterfactual scenario is β_{L_F}, which is the specific price change in France in low competitive areas (less densely populated). Similarly, the coefficients β_{M_B} and β_{H_B} measure Belgian stores' tax pass-through in medium and high competition areas compared to their respective counterfactual in France. The results of this estimation are displayed in Table 3.8 below.

Interestingly, the results of *model 3.7* are similar to those of *model 3.3* and *model 3.5*. Tax shifting decreases with population density (used as proxy for competition). The magnitude of the "competition effect" is also quite similar to the one we find in the previous models. The Wald test on the equality of coefficients indicates that for most products this difference is statistically significant at the 0,01 level (except for product C, where it is significant at the 0,07 level and product F where the competition effect is not significant as in the previous models).

Table 3.8

Population Density as a proxy for Competition (Model 3.7)						
	Product					
	A	**B**	**C**	**D**	**E**	**F**
Low Pop. Density and No Prox. to Lux. (β_{L_B})	**3,48** (0,07)	**2,83** (0,05)	**2,90** (0,08)	**2,89** (0,10)	**2,95** (0,14)	**3,21** (0,14)
Low Pop. Density and Prox. to Lux. (β_{LUX})	**3,37** (0,07)	**2,45** (0,18)	**2,98** (0,08)	**2,79** (0,10)	**2,78** (0,16)	**2,83** (0,22)
Medium Pop. Density (β_{M_B})	**3,27** (0,08)	**2,71** (0,05)	**2,84** (0,14)	**2,65** (0,12)	**2,45** (0,12)	**3,11** (0,07)
High Pop. Density (β_{H_B})	**3,17** (0,06)	**2,47** (0,09)	**2,61** (0,15)	**2,37** (0,12)	**2,28** (0,14)	**3,07** (0,06)
Test on the Equality of Coefficients (H_0: $\beta_{L_B} = \beta_{H_B}$)						
F value	30,55	11,91	3,50	11,05	12,62	1,81
p-value	<0,01	<0,01	0,07	<0,01	<0,01	0,18
Test on the Equality of Coefficients (H_0: $\beta_{L_B} = \beta_{LUX}$)						
F value	12,59	4,63	3,65	3,43	3,57	4,76
p-value	<0,01	0,04	0,06	0,07	0,06	0,03

Notes: All coefficients are statistically significant at the 0,01 level. Standard errors, clustered at the arrondissement level, are in parenthesis. The table displays only the treatment coefficients for Belgium. The 5th and 6th rows show the results of the Wald test on the equality of the coefficients for low and high population density, where the null hypothesis is H_0: $\beta_{L_B} = \beta_{H_B}$. The last two rows show the results of the Wald test on the equality of the coefficients for low density areas close (β_{LUX}) or not close (β_{L_B}) to Luxembourg, where the null hypothesis is H_0: $\beta_{L_B} = \beta_{LUX}$.

As for the "border effect", we find very similar results to *model 3.5* when comparing tax shifting in low competition areas in the proximity or not to Luxembourg. The tax pass-through of stores close to Luxembourg is smaller for most products. The magnitude of these differences is quite similar to the one found in *model 3.5*, with just three of them being significant at the 0.05 level. (i.e., product A, B and F). These results suggest that, after controlling for possible differences in price changes in differently competitive areas in France (by means of population density), the competition effect and the border effect with Luxembourg remain significant.

However, stores in municipalities with different population density might not only differ in terms of the number of local competitors. They could also differ in terms of other demand and supply-side characteristics. For this reason, we re-estimate a different version of *model 3.6*, in which we substitute the number of competing retailers by population density at the municipality level. In this way, we will have a counterfactual for areas with different level of competition (proxied by population density), while also controlling for possible differences in local demand characteristics (proxied by GDP per capita and rural status). The model is as follows:

$$P_{it} = \delta_i + \beta_2 T_t + \beta_3(BE_i \times T_t) + \beta_{Y_F}(T_t \times \ln(Y)_i) + \beta_{Y_B}(BE_i \times T_t \times \ln(Y)_i) + \\ + \beta_{R_F}(T_t \times Rural_i) + \beta_{R_B}(BE_i \times T_t \times Rural_i) + \beta_{D_F}(T_t \times \ln(D)_i) + \\ + \beta_{D_B}(BE_i \times T_t \times \ln(D)_i) + \beta_{LUX}\left(BE_i \times T_t \times LUX_{B_i}\right) + \varepsilon_{it}. \quad (3.8)$$

Where $\ln(D)_i$ is the natural logarithm of population density at the municipality level. The β_{D_F} coefficient measures how price changes in France vary with population density after the reform. This can be interpreted as the counterfactual scenario for Belgian stores facing an increasing competitive pressure at the retail level. The β_{D_B} coefficient measures how the treatment varies with population density, once accounting for this counterfactual scenario. All other variables have the same interpretation as in *model 3.6* and are needed to control for some possible differences in local demand characteristics that can influence tax shifting.

The results of *model 3.8* are displayed in table 3.9 below. For every product, the β_{D_F} coefficient is not significantly different from zero. This indicates that French prices did not change with population density after the reform. The β_{D_B} coefficients are instead all negatives and significant for most products (except for E and F). This suggests a lower tax shifting in more competitive areas (measured by population density). We also re-run the same model by including province-specific treatment effects in order to control for some potential unobservable factors at the province level. However, this does not affect our results.

We run another robustness check in order to verify our results about the border effect with Luxembourg. Although we recognize that this effect is not significant for every product, we would like to verify that the lower tax pass-through for some products in stores close to Luxembourg can be related to cross-border shopping motives. In order to do that do that, we re-estimate a different version of *model 3.5* where we compute the tax pass-through of all stores that are within 50km distance from the Luxembourg border (instead of those within a distance of 10km).[53] The rationale behind this test is to check whether we still find a lower tax pass-through when increasing the distance to the border. If that is the case, then this is somehow concerning as the scope for cross-border shopping should decline with the distance from the

[53] All stores in this area have very few competitors. Therefore, their tax pass-through should tend to be on average larger than in areas with more competing stores.

border, suggesting that perhaps we are probably capturing some other regional effect. The result is that extending the distance to the border to 50 km eliminates the border effect in the sense that we do not find any significant difference in tax shifting between those stores within 50 km from the Luxembourg border and the other stores.

Table 3.9

Population Density as a proxy for Competition with controls for Demand-side Characteristics (Model 3.8)

	Product					
	A	**B**	**C**	**D**	**E**	**F**
"Gross" Treatment	-1,85	4,09*	3,83	-1,33	-6,77*	-2,85*
(β_3)	(1,57)	(2,18)	(5,14)	(3,22)	(3,62)	(2,58)
GDP per capita (FR)	-0,49***	0,04	-0,06	-0,55*	-0,85***	-0,62**
(β_{Y_F})	(0,15)	(0,18)	(0,49)	(0,30)	(0,28)	(0,28)
GDP per capita (BE)	0,60***	-0,01	0,09	0,55*	1,00***	0,64**
(β_{Y_B})	(0,16)	(0,22)	(0,51)	(0,33)	(0,35)	(0,29)
Rural areas (FR)	0,23	0,14	0,44*	0,33*	0,04	0,17
(β_{R_F})	(0,20)	(0,16)	(0,23)	(0,20)	(0,15)	(0,16)
Rural areas (BE)	-0,19	-0,16	-0,58**	-0,47**	0,14	-0,22
(β_{R_B})	(0,11)	(0,16)	(0,25)	(0,22)	(0,20)	(0,16)
Pop. Density (FR)	-0,09	-0,02	0,11	0,05	0,02	0,07
(β_{D_F})	(0,07)	(0,07)	(0,07)	(0,08)	(0,07)	(0,08)
Pop. Density (BE)	**-0,16****	**-0,20****	**-0,29*****	**-0,24****	-0,16	-0,09
(β_{D_B})	(0,07)	(0,08)	(0,09)	(0,10)	(0,10)	(0,08)
Proximity to Luxembourg	**-0,12*****	**-0,44*****	0,03	**-0,15*****	-0,13	**-0,38****
(β_{LUX})	(0,03)	(0,16)	(0,07)	(0,06)	(0,10)	(0,17)

Notes: *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the arrondissement level, are in parenthesis.

3.5.7 Timing of the Tax Pass-Through

So far, we focused on the spatial dimension of the tax pass-through heterogeneity. We have implicitly assumed that the tax shift was uniform over the months after the tax reform. Yet, a tax reform could take some time before being shifted into retail prices and this shift could also vary overtime. Hence, we estimate a model that allows for leads and lags of the treatment effect. On the one hand, this strategy allows us to see how tax pass-through evolved overtime. On the other hand, the leads of the treatment allow testing formally the parallel trend assumption during the months before the tax hike. In particular, these need to be equal to zero, meaning that the spirit price in Belgium and France did not diverge before the tax reform. For each product, we estimate the following model:

$$P_{it} = \delta_i + \sum_{t=-3}^{4} \beta_{F_t} M_t + \sum_{t=-3}^{4} \beta_{B\,t} (BE_i \times M_t) + \varepsilon_{it}. \qquad (3.9)$$

The variable M_t is a dummy variable indicating the month t in which the price is observed. In total, there are eight months in our sample. From August until March. Three months before the tax reform and four months after, plus the month in which the reform is implemented. The month t is indexed such that the month in which the tax reform takes place, which is November, is equal to $t = 0$. In this way, we can refer to t as the number of months before or after the tax reform. We use the month before the tax reform $t = -1$ (October) as the reference month. The coefficients β_{F_t} measure price changes in France over the month before and after the reform with respect to the reference month. All the β_{F_t} with $t \geq 0$ represent the counterfactual scenarios for Belgian stores for each month after tax reform.

The main coefficients of interest in this model are the β_{Bt} coefficients, which measure the price change for each month before or after the tax reform with respect to the reference month (November). Each β_{B_t} with $t < 0$ are the leads of the treatment. In order to see whether the parallel trend assumption holds, these coefficients must be equal to zero. If not, this means that Belgian prices before the tax reform diverged from the French prices and hence we would reject France as being a good control group for Belgium. Yet, our time window before the tax reform is quite narrow, since we can just observe three months before the reform. Each β_{B_t} with $t \geq 0$ measure instead the tax pass-through in the treated group for every month after the tax reform. For instance, β_{B_0} is the tax pass-through during the month of the reform, while β_{B_2} is the tax pass-through two months after the reform. Our empirical test consists in checking whether these effects are statistically different overtime. Table 3.9 shows the results of this estimation.

Although we have already checked for the pre-treatment trend graphically in Section 3.3, the results of *model 3.9* can be quite useful to test the hypothesis of parallel trend before the tax reform. The coefficients measuring the leads of the treatment are not statistically different from zero (with the exception of lead $\beta_{B_{-3}}$ for product E). This indicates that spirit prices in French stores did not diverge from those in Belgium in the three months before the tax reform. The coefficients for the treatment lags indicate that the tax pass-through did generally increase over time after the tax reform. The test on the equality of the tax pass-through one month later and four month later indicates significant difference for four products out of six. Yet, during the first month of tax reform, the tax hike was over-shifted with a confidence level of 95%.

Table 3.10

Time Heterogeneity in the Tax Pass-Through (Model 3.9)

	Product					
	A	**B**	**C**	**D**	**E**	**F**
3 Months Before	**-0,08**	**-0,02**	**0,20***	**0,00**	**0,14****	**0,12**
($\boldsymbol{\beta_{B_{-3}}}$)	(0,06)	(0,07)	(0,11)	(0,06)	(0,07)	(0,14)
2 Months Before	**0,07**	**-0,02**	**0,10**	**-0,06**	**0,09**	**0,09**
($\boldsymbol{\beta_{B_{-2}}}$)	(0,07)	(0,07)	(0,09)	(0,06)	(0,07)	(0,14)
Month of the Reform	**3,00*****	**2,72*****	**2,63*****	**2,63*****	**2,64*****	**3,10*****
($\boldsymbol{\beta_{B_0}}$)	(0,04)	(0,04)	(0,05)	(0,09)	(0,10)	(0,08)
1 Month After	**2,98*****	**2,72*****	**2,91*****	**2,46*****	**2,39*****	**2,84*****
($\boldsymbol{\beta_1}$)	(0,05)	(0,04)	(0,09)	(0,12)	(0,14)	(0,09)
2 Months After	**2,98*****	**2,53*****	**2,89*****	**2,44*****	**2,53*****	**3,10*****
($\boldsymbol{\beta_{B_2}}$)	(0,06)	(0,09)	(0,11)	(0,13)	(0,12)	(0,09)
3 Months After	**3,59*****	**2,61*****	**2,92*****	**2,73*****	**2,66*****	**3,45*****
($\boldsymbol{\beta_{B_3}}$)	(0,09)	(0,10)	(0,10)	(0,12)	(0,13)	(0,11)
4 Months After	**3,69*****	**2,69*****	**3,14*****	**2,87*****	**2,89*****	**3,50*****
($\boldsymbol{\beta_{B_4}}$)	(0,08)	(0,10)	(0,20)	(0,10)	(0,13)	(0,11)
	Test on the Equality of Coefficients ($\boldsymbol{H_0: \beta_{B_0} = \beta_{B_4}}$)					
F value	49,13	0,12	6,03	4,06	3,58	91,17
p-value	<0,01	0,73	0,02	0,05	0,06	<0,01

Notes: *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the arrondissement level, are in parenthesis. The table displays only the betas coefficients for Belgium. The last two rows show the results of the Wald test on the equality of the coefficients for the month of tax reform (β_{B_0}) and 4 months after (β_{B_4}), where the null hypothesis is $H_0: \beta_{B_0} = \beta_{B_4}$. The month before the tax reform $t = -1$ (October) is used as the reference month.

Table 3.11

Short-run vs Long-run Tax Pass-Through Rate

	Product					
	A	**B**	**C**	**D**	**E**	**F**
November	**123%**	**119%**	**115%**	**108%**	**116%**	**128%**
C.I.	121-126	116-123	111-120	101-116	107-125	121-134
March	**152%**	**118%**	**138%**	**118%**	**127%**	**144%**
C.I.	145 -159	110-126	120-155	109-127	115-138	135-153

Notes: C.I. is the 95% confidence interval of the tax pass-through for each product. The tax pass-through is computed with the estimates of *model 3.9*. November is the first month of tax reform. This row shows the tax pass-through in the short-run. March is the last month of price observation. This row shows the tax pass-through in the long-run.

Accounting for timing in tax pass-through can also provide more insights on the competition and the border effects. So far, the analysis of the border and competition effects was carried out by averaging price changes at the store level over the months following the tax reform. The risk is to confound a lower tax shift in more competitive areas with a simple delay in the tax shift needed for those stores to see how competitors react to the reform. The same argument could apply for the border effect, with the stores close to the border waiting to see the effect of the tax reform on cross-border shopping. To test for different timing in the competition and border effect, we estimate a model that accounts for both spatial and time variations in tax shifting. Following *model 3.5,* we specify this model as follows:

$$P_{it} = \delta_i + \sum_{t=-3}^{4} \beta_{F_t} M_t + \sum_{t=-3}^{4} \beta_{L_t} \left(BE_i \times M_t \times Low_{Comp_i} \times NoLUX_{B_i} \right) + \\ + \sum_{t=-3}^{4} \beta_{LUX_t} \left(BE_i \times M_t \times Low_{Comp_i} \times LUX_{B_i} \right) + \sum_{t=-3}^{4} \beta_{M_t} \left(BE_i \times M_t \times Med_{Comp_i} \right) + \\ + \sum_{t=-3}^{4} \beta_{H_t} \left(BE_i \times M_t \times High_{Comp_i} \right) + \varepsilon_{it}. \qquad (3.10)$$

Model 3.10 is a combination of *model 3.5* and *model 3.9.* Each beta coefficient with $t \geq 0$ provides a measure of how tax shifting evolved in areas with different level of competition. This allows us to check whether the "competition effect" on the tax shift is temporary or persistent over the first five months of tax reform. Table 3.12 shows the change in the tax shifting difference between high and low competition areas for each month after the tax reform. The tax shift difference is computed as the difference between the treatment coefficient in high competition areas β_{H_t} and the treatment coefficient in low competition areas β_{L_t}.

As shown in Table 3.12, the tax shifting difference between high and low competition areas becomes statistically significant for all products (except F) two months after the tax reform and it is persistent four months later. The tax shift in high and low competition areas was initially comparable for product B, D and E. Then they start diverging two months later, with the tax shifting in high competition areas being around 0,70€ lower than in low competition areas. This suggests that it took two months before stores adjusted prices in order to account for the competition. For product A and C instead, such difference is already significant during the first month of tax reform. Thus, indicating that prices in low and high competition areas diverged immediately after the tax reform. The results reject the hypothesis that stores facing more competitors tend to delay the tax shift waiting to see how competitors react. Indeed, if that was true we would observe a "front loaded" tax shift difference. Conversely, we find a "back loaded" tax shift difference with the stores in both low and highly competition areas reacting first similarly to the reform and then progressively the competitive pressure introduced some differential adjustment in the tax shifting.

Table 3.12

Timing of the competition effect (model 3.10)

Product	**Competition effect: $(\boldsymbol{\beta_{H_t}} - \boldsymbol{\beta_{L_t}})$**				
	November $(\boldsymbol{t=0})$	***December*** $(\boldsymbol{t=1})$	***January*** $(\boldsymbol{t=2})$	***February*** $(\boldsymbol{t=3})$	***March*** $(\boldsymbol{t=4})$
A	**-0,05**** (0,02)	**-0,12**** (0,05)	**-0,15**** (0,06)	**-0,18***** (0,07)	**-0,06**** (0,03)
B	-0,03 (0,03)	-0,30* (0,15)	**-0,76***** (0,11)	**-0,70***** (0,13)	**-0,74***** (0,14)
C	**-0,31***** (0,05)	**-0,31***** (0,05)	**-0,31***** (0,05)	**-0,30***** (0,05)	**-0,34***** (0,06)
D	-0,10* (0,06)	-0,10* (0,06)	**-0,73***** (0,14)	**-0,72***** (0,14)	**-0,76***** (0,16)
E	-0,04 (0,07)	-0,37* (0,20)	**-0,65***** (0,18)	**-0,37**** (0,18)	**-0,62**** (0,20)
F	0,00 (0,00)	-0,17* (0,10)	-0,02 (0,03)	-0,02 (0,03)	0,01 (0,02)

Notes: The table shows the results of $\beta_{H_t} - \beta_{L_t}$ for each month after the tax reform as estimated in *model 3.10*. The standard errors of this difference are in parenthesis. *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively.

The estimates of *model 3.10* also allow exploring the time dynamics of the tax shifting for stores at the border of Luxembourg. Table 3.13 below displays the timing of the border effect. Interestingly, the table reveals that the border effect on the tax shift appears with some lag (three months after the reform). The tax shift of product B and F was considerably lower in stores close to Luxembourg inducing a price difference between 0,70€ and 1€. The same timing arises for product E but only four months after the reform, with a price difference of 0,78€. Conversely, for product A we find a persistent but negligible difference in tax shifting overtime. These results highlight that it took some time before stores close to Luxembourg adjusted prices differently.[54] A possible explanation could be some demand smoothing during the reform with consumers anticipating the reform by stockpiling spirits just before the tax hike. That is, the

[54] We also estimated a time-varying version of *model 3.4* in order to study the possible timing-varying "border effect" for all the neighboring countries. Yet, we did not find any significant "border effect" apart for Luxembourg.

demand response to the tax hike was postponed for a few months, once the consumers' inventories were over. We confirm the existence of stockpiling in the next section where we study the impact of the tax reform on the quantity of spirits sold in these stores. To check for the robustness of these results, we also estimated *model 3.10* using population density as a proxy for competition (as in *model 3.7*). The results are consistent with the findings of *model 3.10.*

Table 3.13

Timing of the border effect (model 3.10)

Product	**Border effect: $(\boldsymbol{\beta_{L_t} - \beta_{LUX_t}})$**				
	November $(\boldsymbol{t=0})$	***December*** $(\boldsymbol{t=1})$	***January*** $(\boldsymbol{t=2})$	***February*** $(\boldsymbol{t=3})$	***March*** $(\boldsymbol{t=4})$
A	**-0,09***** (0,02)	**-0,07***** (0,02)	**-0,10***** (0,03)	-0,02 (0,02)	**-0,04***** (0,01)
B	-0,03 (0,03)	-0,02 (0,01)	-0,01 (0,03)	**-0,93**** (0,42)	**-0,97**** (0,45)
C	0,02 (0,01)	0,02 (0,01)	0,18*** (0,06)	0,17** (0,07)	0,14* (0,07)
D	0,01 (0,01)	-0,01 (0,01)	-0,36* (0,19)	0,05 (0,05)	-0,02 (0,06)
E	0,04 (0,02)	0,13* (0,07)	0,12* (0,07)	-0,28 (0,21)	**-0,78***** (0,21)
F	0,00 (0,00)	0,27*** (0,07)	**-0,42**** (0,20)	**-0,74**** (0,35)	**-0,81**** (0,39)

Notes: The table shows the results of $\beta_{L_t} - \beta_{LUX_t}$ for each month after the tax reform as estimated in *model 3.10*. The standard errors of this difference are in parenthesis. *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively.

We also use *model 3.10* to test for the parallel pre-treatment trend at the competition subgroup level. The reference month in *model 3.10* is the month before the tax reform ($t = -1$). Each β_{F_t} with $t < -1$ measures the difference in French prices between the reference month and each of its previous month. All other betas with $t < -1$ are the leads of the treatment effect. They indicate how Belgian prices of different competition subgroups differ from this average French price for every month prior the tax reform. To check whether the assumption of parallel pre-treatment trend at the subgroup level holds, it suffices to check whether these leads are not significantly different from zero for every subgroup of stores. This means that the price evolution of these subgroups is parallel to the control group (average French price) and hence parallel to each other. Table 3.A.2 in the appendix shows the leads of the treatment for different degrees of local competition and proximity to Luxembourg. As shown in Table 3.A.2, spirit prices did not diverge in the pre-reform period across different subgroups, with treatment leads being close to zero and not significant.[55]

3.6 The impact on the quantity of spirits sold

In this section, we study the effect of the tax reform on the quantity of spirits sold in the retail chain under consideration. As the tax shifting was substantially heterogeneous over the country, the quantity response to such policy may also vary across store locations. Furthermore, the significant tax shift in areas close to the border also suggests that a great part of domestic sales could have been lost by cross-border shopping. In order to test for these hypotheses, we analyze the number of bottles of spirits that were sold in stores of our retail chain during the period of tax reform. The products we consider are the same six brands analyzed for the tax pass-through estimation. Interestingly, as this retail chain also controls some stores located in the Grand Duchy of Luxembourg, we also have quantity data for stores located on the other side of the border. This allows us to test directly for cross-border sales spillover.

Table 3.14 shows the yearly percentage change in the quantity of bottles sold in each Belgian province (including the Luxembourg province) and in the Grand Duchy of Luxembourg (the country). Overall in Belgium, during the first year of the tax reform (November 2015 – September 2016), spirit sales have declined by 8,51% with respect to the same period in the previous year. Interestingly, sales have continued to drop the year afterwards by 9,25% with respect to the first year of tax reform.[56] The reduction in sales seems quite heterogeneous across provinces. One year after the reform, the sales of spirits in stores located in G.D Luxembourg (the country) have increased by nearly 62% with respect to the previous year. The second year after the reform those sales have continued rising by 72% as compared to the first year of tax reform. These figures suggest massive cross-border shopping of Belgian households in this neighboring country.

To test whether Belgian consumers have anticipated the tax hike by stockpiling spirits, we compare the number of bottles of spirit sold in October 2015 with that of October 2014. The results are shown in Table 3.15 below. Interestingly, we find an increase of nearly 80% in the quantity of spirits sold, which suggests stockpiling in response to the tax announcement in October 2015. If such stockpiling is not properly taken into account when evaluating ex-ante

[55] Few leads are positive for product C and E, but only for the month of August. This can be due to some temporary shock for some stores during that month.

[56] As the tax change was announced in October 2015 (one month before the tax reform), this month is excluded from the computation to remove the possible effect of stockpiling during that period.

the impact of a tax policy, this can lead to overestimating the tax effect on consumer demand (Wang, 2015).

Table 3.14

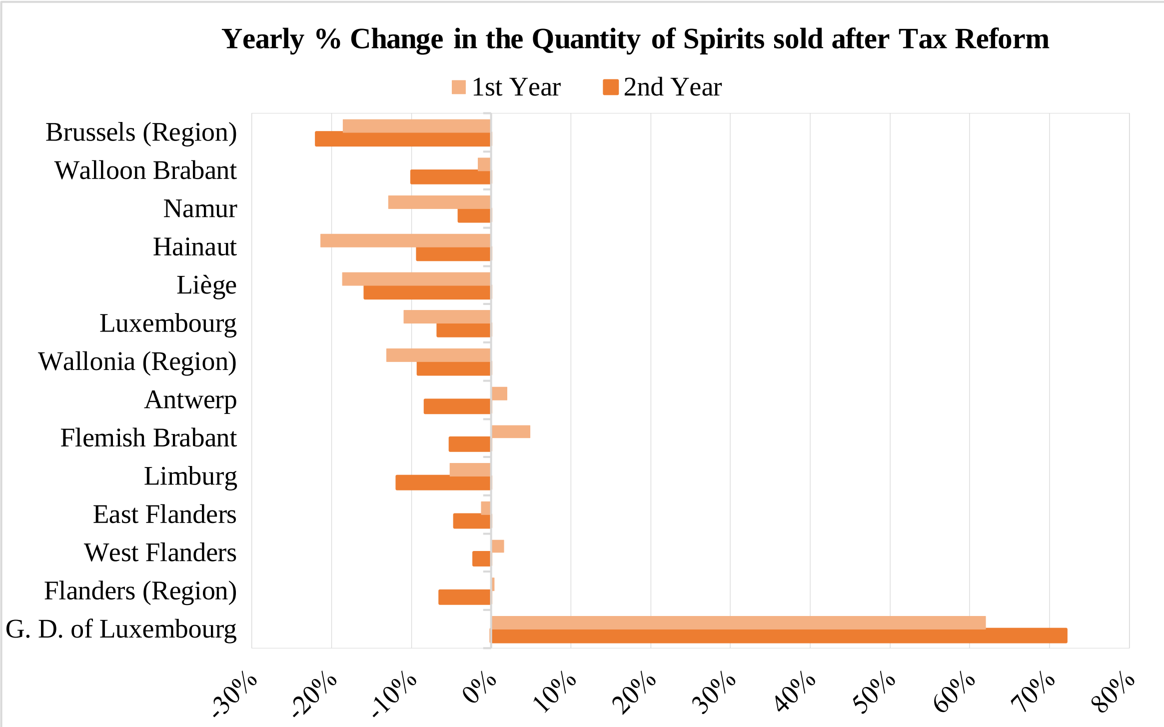

Table 3.15

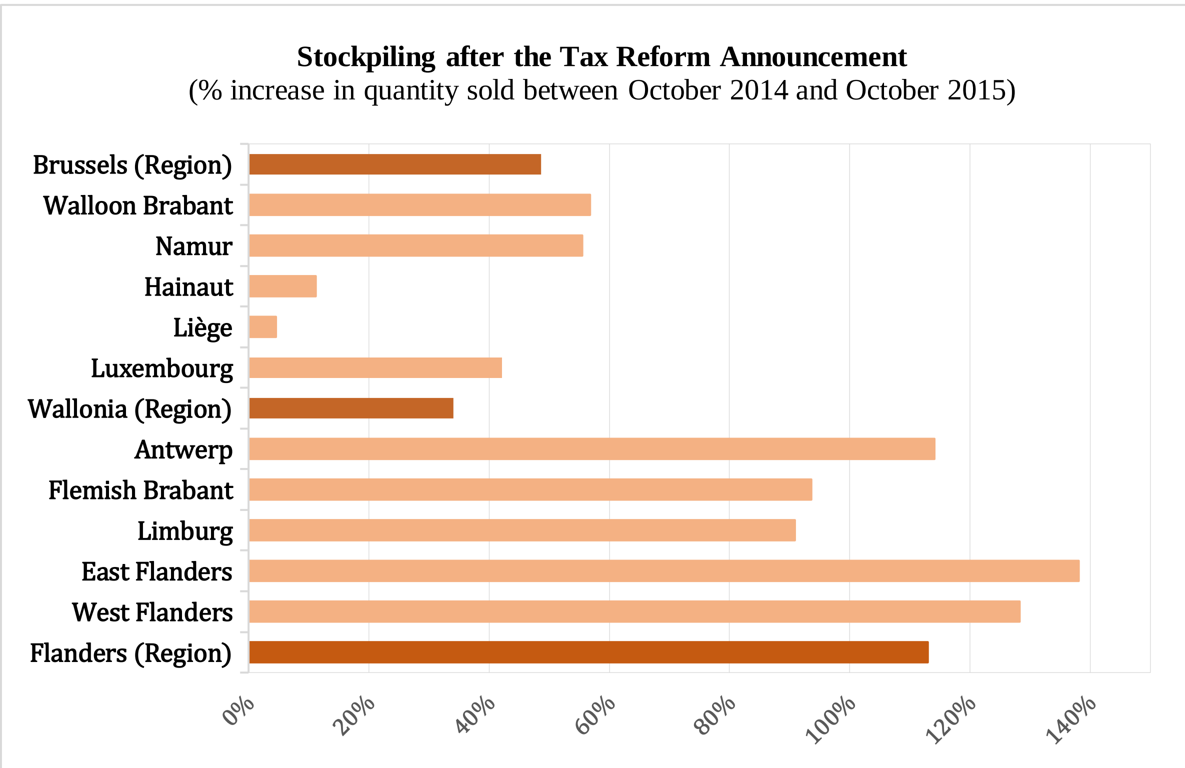

As these figures are limited to only one retail chain, it is not sure whether the tax reform has led some consumer to switch from one retail chain to another. Some evidence of this can be found by looking at the evolution of spirit sales in the provinces of Flemish Brabant, Antwerp and West Flanders during the first year of the reform. In these provinces, stockpiling was greater than average and demand had slightly increased compared to the previous year. Suggesting a possible shift of consumers from other chains and thus an increase in the market share of the chain under consideration. Another possible reason is the lack of alternative as compared to the rest of the country. Indeed, all these provinces are located in the north of the country and share a border with the Netherlands, which is the only neighboring country with similar spirit prices after the tax reform. Conversely, provinces located more in the south (Region of Wallonia), which share borders with countries having lower spirit prices (notably Luxembourg), experienced both a greater drop in demand and a lower spirit stockpiling compared to the average. This can suggest that consumers that have access to cross-border option started purchasing spirits in Luxembourg after the tax reform. Evidence on the evolution of sales in Luxembourg clearly supports this hypothesis.

Since we do not control for any confounding factors that might have occurred during the years after the reform and uses data from just one retail chain of retailers, these figures cannot be interpreted as the causal impact of this tax reform on the volume of sales. Yet, this analysis clearly suggests the presence of stockpiling and the heterogeneous changes in sales across provinces after the tax reform. Moreover, the quantity analysis also reveals a strong positive spillover effect of the tax increase on sales in the neighboring country with the lowest spirit prices (Luxembourg), making the case for cross-border shopping.

3.7 Conclusions

The results of this analysis have shown that the alcohol tax reform implemented in Belgium in November 2015 was mostly over-shifted to the retail price of six major brands of spirit. These products reacted very quickly to the tax reform by adapting their retail prices already during the first month of its implementation. Results also indicate that the tax incidence was substantially heterogeneous both across spirits and over the country. In particular, the intensity of competition is found to be significantly correlated to the extent of tax shifting. The higher the number of retailers in the area, the lower the tax shift. Conversely, proximity to the French, Dutch and German border does not seem to affect the tax shifting even though the tax reform has considerably increased the relative price of Belgian spirits with respect to these countries.

Yet, we do find a quite smaller tax shift for some products in stores close to Luxembourg which is the country having the lowest spirit prices both before and after the tax reform. We have also shown that the tax pass-through varies over time, and that the border and the competition effects are back loaded in the sense that they progressively show up several months after the reform.

In a public health perspective, our findings suggest that the health benefits associated with the tax reform will have a differential impact on Belgian households according to where they live. To support this hypothesis further, we analyze the evolution of spirit sales in the stores considered before and after the reform and provide evidence of a heterogeneous variation of spirit sales over Belgian provinces. We also find evidence of spirit stockpiling before the tax reform and a substantial rise of spirit sales in Luxembourg, which suggests effective cross-border shopping of spirits by Belgian consumers.

References

Ally, A. K., Meng, Y., Chakraborty, R., Dobson, P. W., Seaton, J. S., Holmes, J., ... & Meier, P. S. (2014). Alcohol tax pass-through across the product and price range: do retailers treat cheap alcohol differently? *Addiction*, *109*(12), 1994-2002.

Asplund, M., Friberg, R. & Wilander, F. (2007). Demand and Distance: Evidence on Cross-Border Shopping. *Journal of Public Economics 91*, 141–157.

Bajo-Buenestado, R., &. Borrella-Mas, M.A. (2018). *Passing-through Taxes Beyond Borders with a Cobra Effect*. University of Navarra Working Paper.

Berardi, N., Sevestre, P., Tepaut, M., & Vigneron, A. (2016). The impact of a 'soda tax' on prices: evidence from French micro data. *Applied Economics*, *48*(41), 3976-3994.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, *119*(1), 249-275.

Campos-Vázquez, R. M., & Medina-Cortina, E. M. (2016). *Pass-through and competition: The impact of soft drink taxes using Mexican supermarket stores*. Working Paper.

Carbonnier, C. (2013). Pass-through of per unit and ad valorem consumption taxes: Evidence from alcoholic beverages in France. *The BE Journal of Economic Analysis & Policy*, *13*(2), 837-863.

Cawley, J., & Frisvold, D. E. (2017). The Pass-Through of Taxes on Sugar-Sweetened Beverages to Retail Prices: The Case of Berkeley, California. *Journal of Policy Analysis and Management*, *36*(2), 303-326.

Chandra, A., K., Head, & M., Tappata (2014). The Economics of Cross-Border Travel. *The Review of Economics and Statistics 96*(4), 648-661

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, *99*(4), 1145-77.

Chiou, L., & Muehlegger, E. (2008). Crossing the Line: Direct Estimation of Cross-Border Cigarette Sales and the Effect on Tax Revenues. *B.E. Journal of Economic Analysis and Policy (Contributions)* 8, 1935–1682.

Coibion, O., Gorodnichenko, Y., & Hong, G. H. (2015). The cyclicality of sales, regular and effective prices: Business cycle and policy implications. *American Economic Review*, *105*(3), 993-1029.

Conlon, C. T., & Rao, N. (2016). *Discrete Prices and the Incidence and Efficiency of Excise Taxes*. Working Paper.

DeCicca, P., Kenkel, D., & Liu, F. (2013). Who pays cigarette taxes? The impact of consumer price search. *Review of Economics and Statistics*, *95*(2), 516-529.

DellaVigna, S., & Gentzkow, M. (2017). *Uniform pricing in us retail chains* (No. w23996). National Bureau of Economic Research.

Doyle, J. J., & Samphantharak, K. (2008). $2.00 Gas! Studying the effects of a gas tax moratorium. *Journal of public economics*, *92*(3-4), 869-884.

Einav, L., Leibtag, E., & Nevo, A. (2010). Recording discrepancies in Nielsen Homescan data: Are they present and do they matter? *QME*, *8*(2), 207-239.

Etilé, F., Lecocq, S., & Boizot-Szantai, C. (2018). *The Incidence of Soft-Drink Taxes on Consumer Prices and Welfare: Evidence from the French " Soda Tax".* PSE Working Papers n°2018-24.

Fullerton, D., & Metcalf, G.E., (2002). Tax incidence. In: Auerbach, A., Feldstein, M. (Eds.), Chapter 26 in *Handbook of Public Economics*, vol. 4.Elsevier, Amsterdam, pp. 1788–1872.

Gopinath, G., Gourinchas, P. O., Hsieh, C. T., & Li, N. (2011). International prices, costs, and markup differences. *American Economic Review*, *101*(6), 2450-86.

Griffith, R., O'Connell, M. & Smith, K. (2019), Tax design in the alcohol market, Journal of Public Economics, 172, 20-35.

Grogger, J. (2017). Soda taxes and the prices of sodas and other drinks: evidence from Mexico. *American Journal of Agricultural Economics*, *99*(2), 481-498.

Harding, M., Leibtag, E., & Lovenheim, M. F. (2012). The heterogeneous geographic and socioeconomic incidence of cigarette taxes: Evidence from Nielsen Homescan Data. *American Economic Journal: Economic Policy*, *4*(4), 169-198.

Hindriks, J., & Myles, G. D. (2013). *Intermediate public economics*. MIT Press.

Kenkel, D. S. (2005). Are alcohol tax hikes fully passed through to prices? Evidence from Alaska. *American Economic Review*, *95*(2), 273-277.

Manuszak, M., & Moul, C. (2009). How Far for a Buck? Tax Differences and the Location of Retail Gasoline Activity in Southeast Chicagoland, *The Review of Economics and Statistics 91*, 744–765.

Miravete, E., Seim, K., & Thurk, J. (2017). *One markup to rule them all: taxation by liquor pricing regulation* (No. w24124). Wharton school, University of Pennsylvania.

Nielsen (2017). *Nielsen Grocery Universe Report – Belgium*.

Seade, J. (1985). *Profitable cost increases and the shifting of taxation: equilibrium response of markets in Oligopoly* (No. 260). University of Warwick, Department of Economics.

Shrestha, V., & Markowitz, S. (2016). The Pass-Through of Beer Taxes to Prices: Evidence from State and Federal Tax Changes. *Economic Inquiry*, *54*(4), 1946-1962.

Stolper, S. (2016). *Who bears the burden of energy taxes? The critical role of pass-through. Harvard Environmental Economics Program*, Harvard Kennedy School, Discussion paper 16-70.

Wagenaar, A., Salois, M., & Komro, K. (2008). Effects of beverage alcohol price and tax levels on drinking: a meta-analysis of 1003 estimates from1 112 studies, *Addiction*, *104*, 179–190

Wang, E. Y. (2015). The impact of soda taxes on consumer welfare: implications of storability and taste heterogeneity. *The RAND Journal of Economics*, *46*(2), 409-441.

Weyl, E. G. & Fabinger, M. (2013). Pass-through as an Economic Tool: Principles of Incidence Under Imperfect Competition. *Journal of Political Economy, 121*(3), 528–583.

Xu, X., Malarcher, A., O'halloran, A., & Kruger, J. (2014). Does every US smoker bear the same cigarette tax? *Addiction*, *109*(10), 1741-1749.

3.A Appendix

Figure 3.A.1: Location of French stores (control group)

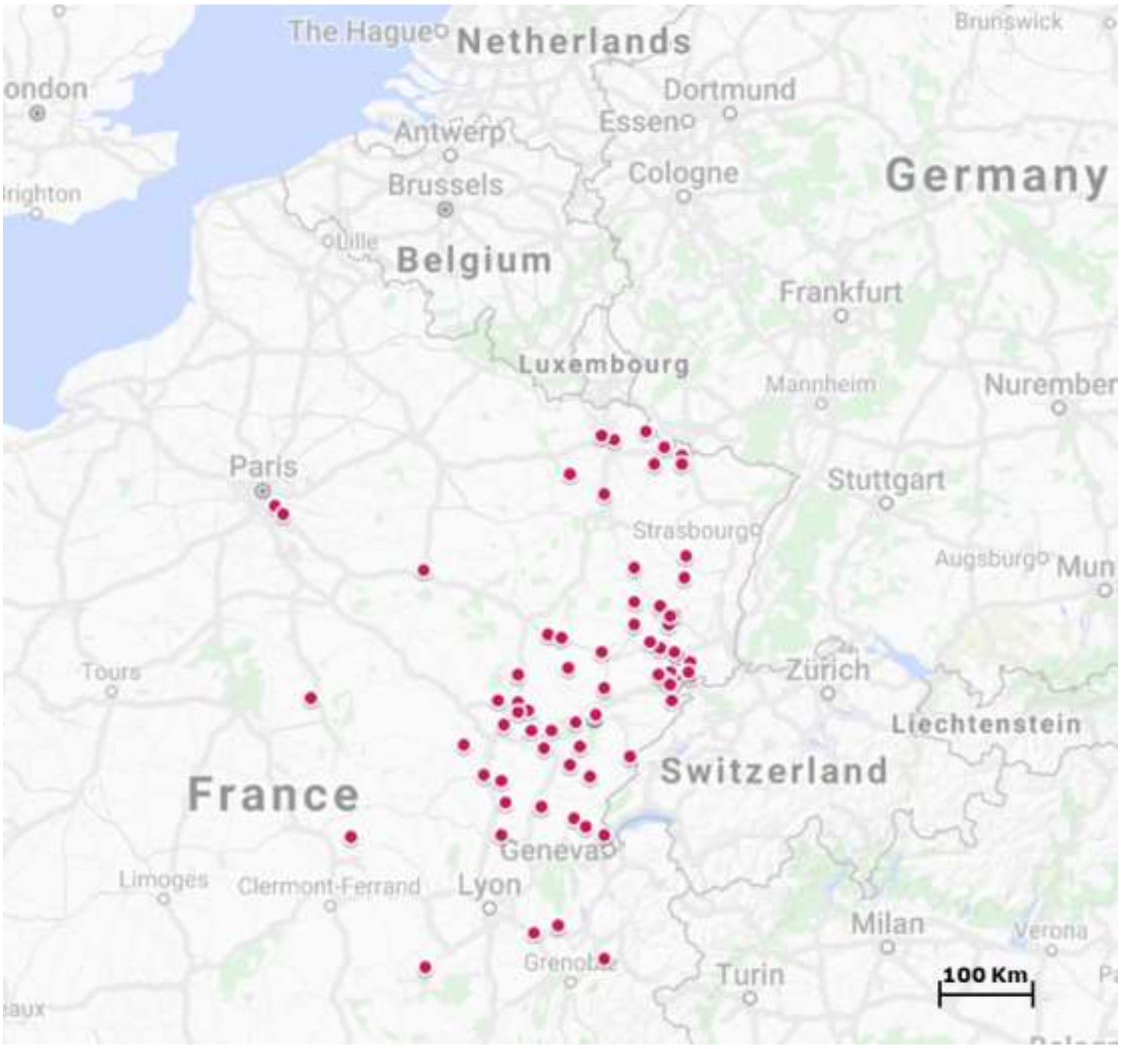

Table 3.A.1

Spatial Price Dispersion						
Product	**Average Price**		**Standard Deviation**		**Levene's Test (homogeneity of σ^2)**	
	Pre-reform	**Post-reform**	**Pre-reform**	**Post-reform**	**F Value**	**P value**
			BELGIUM			
A	15,59	18,79	0,04	0,21	272,06	<0,01
B	11,27	13,86	0,13	0,46	177,12	<0,01
C	10,78	13,44	0,34	0,68	67,45	<0,01
D	13,52	16,27	0,18	0,51	131,42	<0,01
E	15,88	18,36	0,20	0,68	385,33	<0,01
F	15,02	17,98	0,12	0,14	0,70	0,40
			FRANCE			
A	16,31	16,21	0,56	0,50	2,90	0,09
B	11,77	11,68	0,28	0,33	0,01	0,93
C	12,88	12,74	0,51	0,42	1,55	0,22
D	14,36	14,46	0,55	0,48	0,01	0,93
E	15,03	14,97	0,59	0,57	0,53	0,47
F	14,80	14,63	0,64	0,52	4,96	0,03

Notes: *The sample is divided in two groups: Belgium (treated) and France (control). The second column shows the average product price for both groups before and after the tax reform. The third column displays the standard deviation of store prices from the average price before and after the tax reform. The last column shows the results of the Levene's Test on the homogeneity of price variance between the pre and post reform period. The null hypothesis of equal variances between the two periods ($H_0: \sigma^2_{PRE} = \sigma^2_{POST}$) is rejected for all products in the treated group (except for F), while it is accepted for all products in the control group (except for F).*

Table 3.A.2

Pre-treatment trend by subgroups of stores (Model 3.10)

		Product					
		A	**B**	**C**	**D**	**E**	**F**
3 months before reform	**Low Competition** ($\boldsymbol{\beta_{L_{-3}}}$)	-0,08 (0,06)	-0,03 (0,07)	0,06 (0,10)	-0,01 (0,06)	0,14* (0,07)	0,15 (0,14)
	Medium Competition ($\boldsymbol{\beta_{M_{-3}}}$)	-0,07 (0,06)	-0,02 (0,07)	0,07 (0,10)	-0,01 (0,06)	0,17** (0,07)	0,12 (0,14)
	High Competition ($\boldsymbol{\beta_{H_{-3}}}$)	-0,08 (0,06)	-0,01 (0,07)	0,57*** (0,13)	-0,01 (0,06)	0,10* (0,09)	0,09 (0,14)
	Proximity to Luxembourg ($\boldsymbol{\beta_{LUX_{-3}}}$)	-0,08 (0,06)	-0,02 (0,07)	0,13 (0,09)	-0,01 (0,06)	0,18*** (0,09)	0,22 (0,15)
2 months before reform	**Low Competition** ($\boldsymbol{\beta_{L_{-2}}}$)	0,07 (0,07)	-0,03 (0,07)	0,13 (0,09)	-0,06 (0,06)	0,09 (0,07)	0,12 (0,14)
	Medium Competition ($\boldsymbol{\beta_{M_{-2}}}$)	0,07 (0,07)	-0,03 (0,07)	0,12 (0,09)	-0,06 (0,06)	0,09 (0,07)	0,08 (0,14)
	High Competition ($\boldsymbol{\beta_{H_{-2}}}$)	0,07 (0,07)	-0,02 (0,07)	0,04 (0,09)	-0,04 (0,06)	0,07 (0,08)	0,06 (0,14)
	Proximity to Luxembourg ($\boldsymbol{\beta_{LUX_{-2}}}$)	0,07 (0,07)	-0,03 (0,07)	0,13 (0,09)	-0,06 (0,06)	0,10 (0,07)	0,19 (0,15)

Notes: *, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the arrondissement level, are in parenthesis.

Chapter IV

The Incidence of VAT Reforms in Electricity Markets: Evidence from Belgium*

Authors: Jean Hindriks† and Valerio Serse‡

Abstract: In April 2014, the Belgian government reduced the VAT rate on the electricity price from 21% to 6% to support low income families. In September 2015, under the newly elected government, the tax cut was repealed, and the VAT rate was reinstated to 21% in the context of a general tax shift from labour to consumption. This paper investigates the impact of such temporary and exogenous VAT reform on the Belgian electricity market. We study both the pass-through of the VAT reform to electricity prices and the effect of this (exogenous) price change on electricity demand. We estimate the VAT pass-through on residential electricity price by a *difference-in-differences* approach, using business electricity prices (not subject to VAT) as a control group. To estimate the impact of the VAT change on demand, we perform an event study on the electricity flowed monthly over the grid at the network operator level. Our findings reveal that both the tax cut and the tax hike were entirely shifted to the electricity price (i.e., 100%). Exploiting different sources of price variation, our results reveal a price elasticity of residential demand for electricity between -0.09 and -0.17. Interestingly, we also find that demand reacted quickly and symmetrically to the VAT cut and the subsequent VAT hike.

JEL No: H21, H22, H23, Q41, Q48.

Keywords: tax incidence, VAT reform, demand elasticity, electricity markets.

---

* We are grateful to the CREG for the support on data collection and in particular to Christine Cobut, Natalie Cornelis and Laurent Jacquet. We have benefited from very useful comments and suggestions by Paul Belleflamme, Olivier De Groote, Axel Gautier, Anthony Papavasiliou, William Parienté, and various seminar participants at the Economics School of Louvain and Toulouse School of Economics. The financial support from the FNRS is gratefully acknowledged.

† CORE, Université catholique de Louvain, e-mail: jean.hindriks@uclouvain.be.

‡ CORE, Université catholique de Louvain, FNRS-FRESH grant holder, e-mail: valerio.serse@uclouvain.be.

4.1 Introduction

In April 2014, the Belgian government reduced the value-added tax (VAT) rate on the electricity price from 21% to 6%. In September 2015, this tax cut was repealed, and the VAT rate reverted to 21%. This paper investigates the impact of such temporary VAT reform on the Belgian electricity market. The aim is to study the effect of this tax reform on electricity price and demand. The results of this analysis shed light on the tax incidence in the electricity market, and the price elasticity of electricity demand. The tax reform and resulting price change provide an excellent natural experiment, as they are exogenous to households' consumption behavior. We also provide novel evidence on the possible asymmetry in tax incidence following the tax reform. Indeed, the VAT reforms involved first a tax cut to 6%, followed by a tax hike to 21%. This provides a perfect setting to test for symmetry in the price and demand responses to tax cuts and tax hikes within the time interval April 2014-September 2015.

The quality of the tax incidence analysis depends on the exogeneity of the tax reforms. In a Royal Decree of 21 March 2014, it was decided to cut the VAT rate on electricity supply to residential customers from 21% to 6% conditional on the evaluation of this tax reform by September 2015. The reform was part of the "competitiveness and employment pact" of the government Di-Rupo. The VAT cut was in fact a social measure from a socialist government intended to benefit the low-income households, and its budgetary cost was perceived as limited, given its impact on wage and social security benefit indexation. Indeed, the VAT cut was expected to lower the inflation rate, which given the Belgian automatic indexation, translates into a reduction of the wage costs and social security benefits. In October 2014, the center-right government Michel replaced the socialist government. The government Michel enacted the Plan Bureau to estimate the budgetary cost of the reform. The inflation rate was reduced by 0.4 pp, indexation was postponed by four months, and the net budgetary cost of the reform was estimated at 536 million euros in 2015 (BFP, 2015).

Then the government Michel proposed in July 2015 a new "tax shift" plan. This plan organized a broad shifting of tax base from labor to consumption as part of a strategy to lower wage costs and promote employment. The tax shift plan involved the VAT rate increase from 6% to 21% in September 2015. The alcohol tax reform studied in chapter 3 was also part of the same tax shift plan. The VAT reinstatement was indeed a political move from a liberal coalition to undo the reform from the socialist government. Therefore, the two successive VAT reforms can be seen as exogenous shocks motivated by political decisions from two different governments, with no interaction with the electricity market.

The structure of the electricity market is rather specific. The retail market is competitive, whereas transmission and distribution markets are natural monopolies regulated by the Belgian energy authority, the Commission for Electricity and Gas Regulation (CREG). This means that energy producers and service providers (retailers) can react freely to a variation in VAT. Firstly, we study whether these operators entirely shifted this temporary VAT cut to the electricity price paid by households or retained a part of it in their price. We estimate the VAT pass-through to the residential electricity price by employing a *difference-in-differences* strategy. We use as a control group the electricity price paid by professional consumers (firms) to assess the impact of the temporary VAT reform of the electricity price paid by households. The assumption is that the electricity price for households would have followed the same trend as the electricity price for firms in case the reform did not occur. The VAT is charged only on final consumption by households but not on intermediate consumption by firms and professional consumers.

The results show that the VAT cut was entirely shifted to the household electricity price. Producers and service providers passed the tax cut entirely to the benefit of households, with both regulated and unregulated firms reacting similarly. Furthermore, we do not find any asymmetry in tax shifting between the first VAT cut to 6% and the subsequent VAT hike to 21%. These findings highlight that the tax change is entirely shifted to electricity prices. If this result is typical for the regulated component of the electricity bill (transmission and distribution), it is less so for the unregulated components (service providers and energy producers). Given that tax changes are perfectly shifted on the residential electricity price, these can have a significant impact on both electricity demand and tax revenues.

We then estimate the impact of the VAT reform on residential demand for electricity. We study the monthly demand of electricity at the network distributor level, controlling for changes in other determinants of energy use during the period. We find that the VAT cut from 21% to 6% generated around a 2% increase in demand for electricity. Exploiting different sources of price variation, our results reveal a price elasticity of residential demand for electricity between -0.09 and -0.17. This measure, however, can also vary across regions. We provide novel evidence about the symmetry in the demand response to price cuts and price hikes. Demand changes to the same extent (albeit in a different direction) after a price cut and a price hike of the same magnitude. This finding has important implications for price-based climate policies, as it indicates that price hikes through taxation are entirely shifted to consumers and can be effective in reducing energy use.

Lastly, we show that consumers responded quickly to the VAT reform, increasing their demand one month after the VAT cut. This evidence contrasts with the existing literature, which suggests that consumers respond slowly to changes in electricity prices.[57] The VAT reform was announced early in the media, and the pass-through was perfect, making the reform more salient. Such salience of the tax reform may induce consumers to react more quickly compared to other price changes (Chetty, Looney, & Kroft, 2009). We also find that the increase in electricity demand is mostly concentrated in sunnier and warmer periods when energy is less needed for both heating and home lighting. This result can cast doubts on the welfare effect of such a policy, given that electricity production involves negative externalities from higher CO2 emissions.[58]

The paper is organized as follows. In section 2, we review the existing literature and identify our contribution. Section 3 describes the structure of the Belgian electricity market. Section 4 presents some theoretical background on VAT incidence and market structure. In section 5, we describe the data used. Section 6 presents the models and shows the results of the empirical analysis. Section 7 shows the results of some robustness checks. Section 8 concludes.

4.2 Background Literature

This work contributes to the literature on VAT incidence by providing novel insights in the context of electricity markets. The empirical research has mostly focused on the incidence of excise taxes rather than sales taxes and VAT.[59] Recent studies have found that the cost of emission permits in Europe is almost entirely shifted to electricity prices (Fabra & Reguant, 2014; Hintermann, 2016). However, there is a lack of evidence regarding the pass-through of value-added taxes to retail electricity prices. Most studies estimate sales taxes (or VAT) pass-through following a reduced-form approach between prices and tax rates (see, for instance, Poterba, 1996; Besley & Rosen, 1999). Doyle and Samphantharak (2008) and Marion & Muehlegger (2011) study the incidence of sales taxes on fuel in the United States. They both find a pass-through close to unity, which can vary with differences in supply-side conditions.

[57] Price elasticity of electricity demand tends to increase over time because energy consumption depends on the stock of the owned appliances. While this is fixed in the short-run, this can be changed in the long-run as a function of the present and expected future prices.

[58] Around the 30% of the electricity generation in Belgium is based on fossil fuels, mainly natural gas (Eurostat, 2020).

[59] Many works studied excise tax incidence in the context of sodas (Berardi et al., 2016; Cawley & Frisvold, 2017; Etilé, Lecocq, & Boizot-Szantai, 2019), cigarettes (Harding, Leibtag, & Lovenheim, 2012; DeCicca, Kenkel, & Liu, 2013) and alcoholic beverages (Kenkel, 2005; Shrestha & Markowitz, 2016; Hindriks & Serse, 2019).

Doyle & Samphantharak (2008) also provide some evidence that tax cuts are less than fully shifted to fuel prices.

Carbonnier (2007) studies the VAT incidence of two French reforms decreasing the VAT rate on housing repairing services and new car sales. He finds that both VAT cuts were not entirely shifted to consumer prices. The tax shifting was 77% for housing repairing services and 53% for new car sales. He associates such a difference in the VAT pass-through to the higher concentration in the new car sales sector. Kosonen (2015) studies the incidence of the VAT reform on hairdressing services in Finland, which reduced the VAT rate from 22% to 8%. He finds that hairdressers have only passed half of the VAT reduction to consumer prices. Benedek et al. (2019) review the literature on VAT pass-through in the Eurozone for a wide variety of markets, but they do not consider the one for electricity. They find that the pass-through of the standard VAT rate is usually 100%, while for reduced rates, it is around 30%. Contrary to these findings, we show that the reduction from a standard rate of 21% to a reduced rate of 6% was entirely shifted to retail electricity prices.

This literature typically identifies the VAT pass-through by looking at the change in consumer price before and after the VAT reform, while controlling for the price evolution in other markets.[60] Yet, finding a reliable control group to VAT changes is rather difficult. This is because the VAT affects the whole market to which it is applied. Furthermore, even if the VAT change affects only one sector, there is a risk of a cross-price effect on the closest sectors producing substitute products. We extend this literature by adopting a *difference-in-differences* approach that uses business prices as a counterfactual scenario. Business prices are not subject to VAT but share the same cost components with residential prices. Moreover, these two tariffs are not substitutable since households cannot get a business price as residential customers. Our findings suggest that without controlling for business prices, we would have reached the misleading conclusion that the VAT cut was not entirely shifted to consumer prices.

The literature focussing on the price elasticity of electricity demand is instead more substantial, although very few papers are based on *quasi-experimental* (exogenous) variation in prices. Most studies estimate dynamic panel models on state-level data. Labandeira, Labeaga, & López-Otero (2017) provide an exhaustive review of this literature, showing that the short-run and the long-run demand elasticities are on average -0.21 and -0.67, respectively. However, a

[60] Benedek et al. (2019) use as counterfactual price the one of the same commodity sold in other European countries that did not change their VAT rate. Carbonnier (2007) uses the price index for non-treated industries. While Kosonen (2015) uses the price of other labor intensive services.

limitation of these studies is that they tend to assume that prices are exogenous to consumer demand. Furthermore, it has been shown that estimates of dynamic panel models can vary significantly with the specification used (Alberini & Filippini, 2011).

Some recent studies exploit *quasi-experimental* variation in prices to recover the price elasticity of electricity demand. However, those studies do not consider tax change. Ito (2014) estimates the demand response to both average and marginal electricity price variations induced by a discontinuity in electricity service areas in California. He finds that the demand is sensitive to average prices rather than marginal prices, with a demand elasticity of -0.09 four months after the (average) price changes. In our context, we do not disentangle average and marginal prices as the VAT is applied to both the price per kWh and the fixed fees. Therefore, the VAT reform changes marginal and average prices to the same extent.

Deryugina, MacKay, & Reif (2020) study the elasticity of residential electricity demand following an exogenous drop in electricity prices due to a municipal aggregation policy in Illinois. The municipal aggregation allowed local communities to select new electricity suppliers on behalf of their residents with the approval of a local referendum. This policy led to a large price decrease in communities that implemented aggregation. They find a demand elasticity that can vary between -0.09 and -0.16 within 6 and 12 months, respectively. Nevertheless, as they only observe a price cut, they cannot check for differences in demand elasticity following a price hike.

Interestingly, our measure for the price elasticity of demand is in keeping with these quasi-experimental studies. In contrast with these studies, however, we find that the demand response to the price change is immediate. This is probably because we are exploiting a price variation that was less complex and more salient for consumers. Given that we do not have a reliable control group for the demand analysis, we run a *placebo test* on the evolution of the electricity demand of SMEs (not subject to VAT). We find that during the VAT cut, the SMEs did not increase electricity demand, contrary to households. Therefore, this is suggesting that our estimates are likely to capture the causal effect of price on demand.

4.3 The Belgian Electricity Market

Electricity markets have a unique structure. Electricity is firstly produced by electricity generators, and it is then transported to end-users. The transport is composed of transmission and distribution networks. Transmission system operators manage the transmission network. They are in charge of transporting electricity over long distances from the generators to the distribution network. The distribution network is instead managed by distribution operators, which deliver electricity to the end-users. The electricity is sold to consumers by service providers, who buy it from the generators or produce it themselves. The electricity bill paid by consumers is the sum of the costs of these different operators. As for most countries, electricity transmission in Belgium has regulated tariffs. This means that any change in price for the transmission and distribution of electricity has to be approved ex-ante by the energy regulator (CREG). Downstream, service providers also set their prices freely for the supply of electricity in a competitive market.

Upstream, the energy component of the electricity bill is set freely by producers since the liberalization of the electricity and gas market. This price includes the pure energy component and the green energy component (22.1% of the total energy component in 2019 or 7.8% of the total cost of electricity). Between 2013-2017, the government used a "safety net" mechanism to limit the price of the energy component upward. Under this mechanism, energy suppliers can index variable energy prices every three months. The CREG monitors indexing and compares possible price increases with those of consumers of the same type in neighboring countries.

Another important component of the electricity bill is the taxes levied at both federal and regional levels. The shares of the various components may vary depending on the type of customer considered, distribution areas, regions, and suppliers. In this work, we consider the marginal price of electricity. Hence, we consider only those taxes levied per kWh consumed. These include both regional and federal levies, which are also subject to VAT.

4.4 VAT Incidence and Market Structure

Electricity tax is levied on the final supply of electricity to the consumer, and the liability arises at the time the electricity is supplied. The supplier is responsible for payment of the tax and all remunerations. In this paper, we consider the tax incidence analysis of the VAT change for households only (domestic electricity use by contrast to business use). Before we carry out the empirical analysis, let us remind some theoretical features of tax incidence depending on the market structure and the type of taxes (ad-valorem versus specific taxes). Specific tax is a tax per unit (kWh) for electricity supplied. The ad valorem tax, such as the VAT, is a tax per unit of the price (€ cents) for electricity supplied.

The two types of taxation can be represented as follows. Let q denotes the consumer price and p the producer price. A specific tax rate t induces a consumer price of $q = p + t$. An ad valorem tax rate of τ induces a consumer price of $q = (1 + \tau)p$. Tax incidence measures the impact of a change in the tax rate on the consumer price. Fixing the ad valorem tax rate at $\tau = t/p$ implies that $q = p + t = (1+\tau)p$ which induces the same consumer price q and, therefore, the same consumption x. Both taxes would then raise the same tax revenue. Indeed under ad valorem tax, the revenue raised is $R = \tau p x$ which using $\tau = t/p$ is equivalent to $R = tx$, which is the revenue raised under the specific tax. The market equilibrium determines the incidence of the tax on producer and consumer prices by equating demand and supply $D(p) = S(q) = x$.

Figure 4.1: Tax incidence under perfect competition

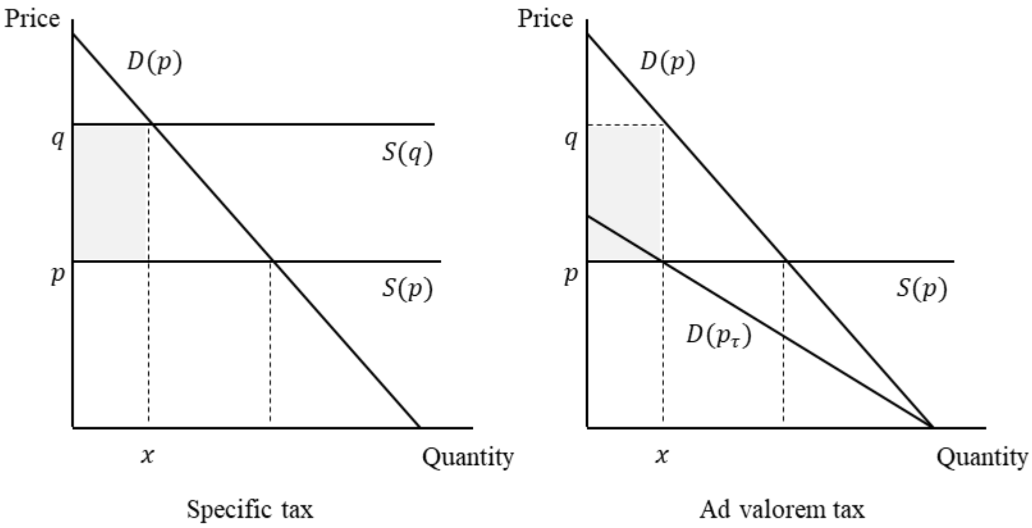

(Tax revenue in shaded area)

Under imperfect competition, the equivalence breaks down, and ad valorem tax dominates specific tax in the sense that ad valorem tax raises more revenue for the same consumer price than the specific tax. The reason is that under imperfect competition, the producers are no longer price takers, but they can raise the price above the competitive price by limiting the supply. By taxing price directly rather than quantity, ad valorem taxation will limit this market power. To see this, let us consider a monopole facing a demand $x(q)$ with a constant unit cost c. The monopoly profit under specific tax rate t is $\prod = (q-t)x - cx = qx - (c+t)x$. The specific tax is equivalent to increasing the unit cost at a rate t. The ad valorem tax at rate τ implies that for consumer price q, the monopoly receives $p = q/1+\tau$. From the monopoly's perspective, the specific tax increases cost from c to $c+t$, whereas the ad valorem tax reduces revenue from qx to $qx/1+\tau$.

Under perfect competition, the supplier perceives no effect on price, and so both taxes are equivalent. However, under imperfect competition, it is no longer valid. With ad valorem taxation, monopoly profit $is \prod = qx/1+\tau - cx = [1/1+\tau]\,[qx-(c+\tau c)x]$. The ad valorem tax is equivalent to combining a profit tax at rate $\tau/1+\tau$ and a specific tax at rate τc. The profit tax is proportional to profit, and so it does not affect supply. If we set ad valorem tax rate such that $\tau c = t$ the ad valorem tax induces the same consumer price q as the specific tax t. But the ad valorem tax raises more revenue than the specific tax since $R = \tau c x + (\tau/1+\tau)\prod = tx + (t/c+t)\prod > tx$. Thus, the ad valorem tax raises more revenue than the specific tax for the same consumer price. Alternatively, for the same tax revenue, ad valorem tax reduces consumer price compared to a specific tax. Ad valorem tax is superior because it is a tax on the market power (profit margin) of the supplier.

4.5 The Data

The price data used in this study are provided by the Belgian Energy Regulator (CREG). We build a price series for every type of electricity contract. Any service provider can offer electricity contracts to either firms or households. Only households are subject to VAT since firms can fully recover the VAT paid from the VAT received on sales. Most electricity providers in Belgium have different tariffs for different consumption profiles. In this study, we use the price of electricity paid by households with a consumption profile of 3,500 kWh. For firms, we consider the tariff for a consumption profile of 50,000 kWh. We use data from two leading electricity providers in Belgium: Luminus and Electrabel (Engie). We use data from their two most common contracts. These are *Luminus Click*, *Luminus A+*, *Easy Fixed*, and *Easy Indexed.* Each of these contracts has both the firm and household version. Hence, we consider eight price series of electricity contracts, where prices are expressed in €cent/kWh.

The total price paid by the client for a given electricity service is the sum of various cost components. These components include the cost of energy, distribution & transmission, and taxes & surcharges. Service providers are responsible for the energy supply to the client and bill the energy use. The price for the contracts *Luminus Click* and *Easy Fixed* is fixed overtime for a client. This means that once the client signs these contracts, the electricity price for that client is fixed for a given period (this can be from 1 up to 3 years). That does not prevent the client to switch (without cost) to a better contract offered by a different service provider. Price comparators are publicly available. Even though these prices are fixed for the client who already signed it, they can change monthly for any potential new client (Newcomer contract). The other contracts are instead indexed. Their price can change for both new and old customers according to an indexation formula. The price, however, can be indexed just every quarter, which means that these prices only change three times per year.

Another price component is the remuneration for the network operators, which are responsible for the distribution and the transmission of the electricity in the grid. These operators are active locally, and their prices can vary over space. In this study, we have data for a total of 23 distribution grid operators.[61] Although these fees can vary across localities, they do not vary across different service providers and types of electricity contracts. That is, every network operator has to charge the same price to every provider locally. The rest of the bill is composed of federal and local energy taxes. If the contract is designed for a household, the VAT is then applied to each price component. The total monthly electricity price for a client is the sum of all these tax inclusive price components. This price can vary across both contracts offered by service providers and local network operators. Table 4.1 below summarizes the data in our sample, distinguishing between the electricity prices for households and firms during the period of VAT reforms.

Table 4.1

Summary Statistics – Electricity Price

	21% VAT period (06/13 – 03/14)		**6% VAT period (04/14 – 08/15)**		**21% VAT period (09/15 – 03/16)**	
Price Component	*Households*	*Firms*	*Households*	*Firms*	*Households*	*Firms*
Energy	**7.41**	6.53	**6.45**	6.49	**6.71**	5.94
Green Quota[62]	**1.84**	1.52	**1.84**	1.74	**2.46**	2.03
Distribution	**9.92**	8.25	**9.09**	8.61	**11.77**	9.74
Transmission	**2.40**	1.96	**2.49**	2.34	**2.87**	2.37
Other Taxes (VAT excl.)	**0.64**	0.53	**0.51**	0.50	**0.53**	0.51
Total	**22.21**	18.80	**20.38**	19.68	**24.35**	20.59

Notes: *Price components are expressed in €cent/kWh. These prices are the average over the contracts and network operators included in the sample. For Household prices, each component includes the VAT.*

The summary statistics reveal a significant drop in the electricity price over the period of the temporary VAT cut at 6%. However, the costs of the different components may also have changed over that period, sometimes under the influence of different policy changes. To

[61] However, in the consumption analysis we drop observation about eight of them, as we do not have enough data about consumption for these distributors before the VAT reform. These distributors are also small ones.

[62] The green quota refers to a contribution for the production of renewable energy, which firms pass to consumers. Firms are obliged to purchase a minimum amount of green certificates to support the development of renewable energy sources.

identify the impact of the VAT change on prices, we should carefully consider the counterfactual evolution of electricity prices in the absence of the VAT change. We will do that using the change of the price for firms as a counterfactual.

Turning to consumption data, they are provided by *Synergrid*, the Belgian federation of network operators. The consumption data consists of the monthly amount of kWh flowed over the electricity grid through each distribution system operator (DSO). Hence, the unit of observation here is the monthly electricity consumption at the DSO level. We have information about 15 DSOs over a period of four years (January 2013 – December 2016), which makes a total of 720 observations. This measure does not account for the electricity consumed by households producing their own electricity (e.g., through solar panels). Although the auto production of electricity is becoming increasingly popular in Belgium (due to generous subsidy schemes[63]), the majority of electricity consumed by Belgian households is observable since it flows over the grid. For the demand analysis, we match the price data described above for every DSO. Regulated tariffs and surcharges are specific to the given DSO. However, as energy costs can vary across service providers, we use the average energy tariffs across service providers in our sample.

Our data can distinguish electricity demand for different consumption profiles. Each DSO has information about the flow of electricity passing through its network every quarter of hour thanks to reding meters located at the main injection points. This flow of electricity is tracked for 4 different consumption profiles: businesses with a power supply <56 kVA; businesses with a power supply ≥ 56 kVA; residential with a usage ratio night/day < 1.3; residential with a usage ratio nigh/day ≥ 1.3. This distinction is possible because each of these profiles has a different electrical connection to the grid, which is chosen based on the client needs. For household demand, we use the data of households having a usage ratio night/day below 1.3.[64] This consumption profile tends to have a yearly consumption close to 3,500 kWh, which corresponds to the price series we consider and to the most common consumption profile. The consumption data also allows separating the consumption of households benefiting from social tariffs. This is possible as these households are furnished by a separate social supplier. A distinction that is important as these households face different prices.[65]

[63] See Boccard & Gautier (2020) and Gautier & Jacqmin (2020) on the impact of subsidies for solar panels on PV installations in the Belgian region of Wallonia.

[64] This correspond to the Synthetic Load Profile S21 provided by *Synergrid*.

[65] The price series for social tariffs is displayed in Figure 4.A.1 in the appendix. These tariffs are fully regulated and tend to vary much less than normal prices.

We also collect data about other factors that could have affected the evolution of electricity demand. These include the hours of sunlight, the degree-day,[66] and the amount of electricity that is auto produced (i.e., the electricity not flowing over the grid). Table 4.2 below shows an overall decrease in electricity demand over time. However, various factors affecting electricity demand did vary as well over the same period. The degree-day tends to decrease over the three periods, meaning that there was less need for energy to heat buildings given the higher outside temperature. The level of sunlight increased during the last period, suggesting that less electricity for house lighting was needed. The level of auto production of electricity increases from 2014, which could partially explain the decrease in electricity demand in the network grid over the last two periods (from April 2014 to December 2016).

Table 4.2

Summary Statistics – Electricity Demand

Variable	21% VAT period (01/13 – 03/14)	6% VAT period (04/14 – 08/15)	21% VAT period (09/15 – 12/16)
Consumption (kWh)	**267.81**	**255.90**	**245.39**
Auto production (TWh)	**1.46**	**1.80**	**1.75**
Sunlight	**122.42**	**122.19**	**156.33**
Degree-day	**224.20**	**188.29**	**142.94**

Notes: *Consumption is the average monthly consumption per EAN code in kWh. The hours of sunlight and degree-day are the monthly average over the mentioned period. Auto production is a yearly estimate at the national level and is expressed in TWh (it is not seasonally adjusted).*

4.6 Empirical Strategy

We estimate the impact of the temporary VAT reform on both the price and demand for electricity. The pass-through of this VAT change on the residential electricity price is calculated through a *difference-in-differences* analysis. While we estimate the impact on demand by means of an event study regression, where we control for other possible factors affecting the electricity use over the same period. To assess the VAT incidence on residential electricity prices, we use as a control group the price of electricity paid by firms. Hence, we use the electricity price for firms as the counterfactual for the evolution of the electricity price for the household if the VAT

[66] The degree-day is a measure of the amount of energy needed during a day to heat an average building. This data is also taken from *Synergrid*.

reform was not implemented. The reason is that households are subject to the VAT but not the firms (who can recover the VAT).

The electricity price for firms is a reliable control group for this analysis. The two groups share the same cost components and are mostly provided by the same electricity providers. The main difference in the final price per kWh is the VAT paid by households. Figure 4.3 shows the evolution of the electricity price for both firms and households. The graph suggests that both treated (household) and control (firms) have a parallel trend both before and after the VAT cut, apart for the two periods of VAT change. Therefore, this indicates that the electricity price of firms is a reliable control group for the price households would have faced in the absence of VAT change.

Although household and firm tariffs mostly share the same cost components, the same cost shock may not be passed through equally. This is because household tariffs are subject to VAT, while firm tariffs are not. Therefore, any cost shock may be passed through to household prices differently than to firm prices. For instance, a 1 €cent/kWh increase in the cost of producing energy is likely to increase firm prices by 1 €cent/kWh, while household prices should rise by 1 €cent/kWh plus VAT (with 21% VAT this is 1.21 €cent/kWh).

Figure 4.2: Cost pass-through with VAT

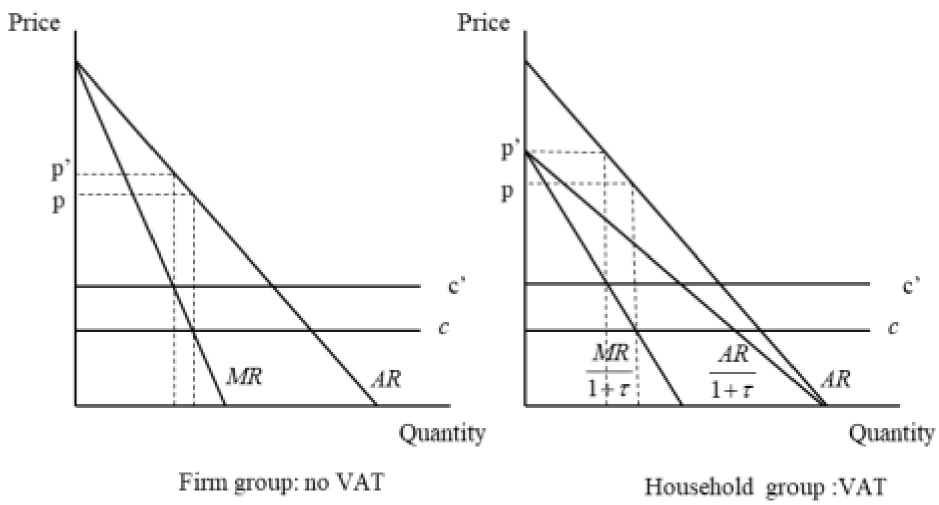

This difference in cost shifting is displayed in Figure 4.2 above. The figure shows how the same increase in marginal cost is shifted to prices with and without VAT under imperfect competition. The left-figure shows the increase in the electricity price from p to p’ for firms (not subject to VAT) following a marginal cost increase from c to c’. The right-figure instead shows how the same cost shock affects the equilibrium price for households (subject to VAT).

The difference in cost shifting is because the VAT reduces both the marginal (MR) and average (AR) revenues of electricity suppliers. Hence, leading to a lower quantity and a higher price after the cost increase.

We adjust our control group to account for this asymmetry in cost shifting by including the VAT to the electricity price for firms. As this is the counterfactual scenario for household prices in the absence of VAT change, we added a 21% VAT (equal to the tax rate before the tax change) to firm prices for the entire period. In such a way, the control group is the electricity price for firms assuming they were subject to the pre-reform VAT rate of 21%. This price series is displayed in figure 4.3 below.

Figure 4.3

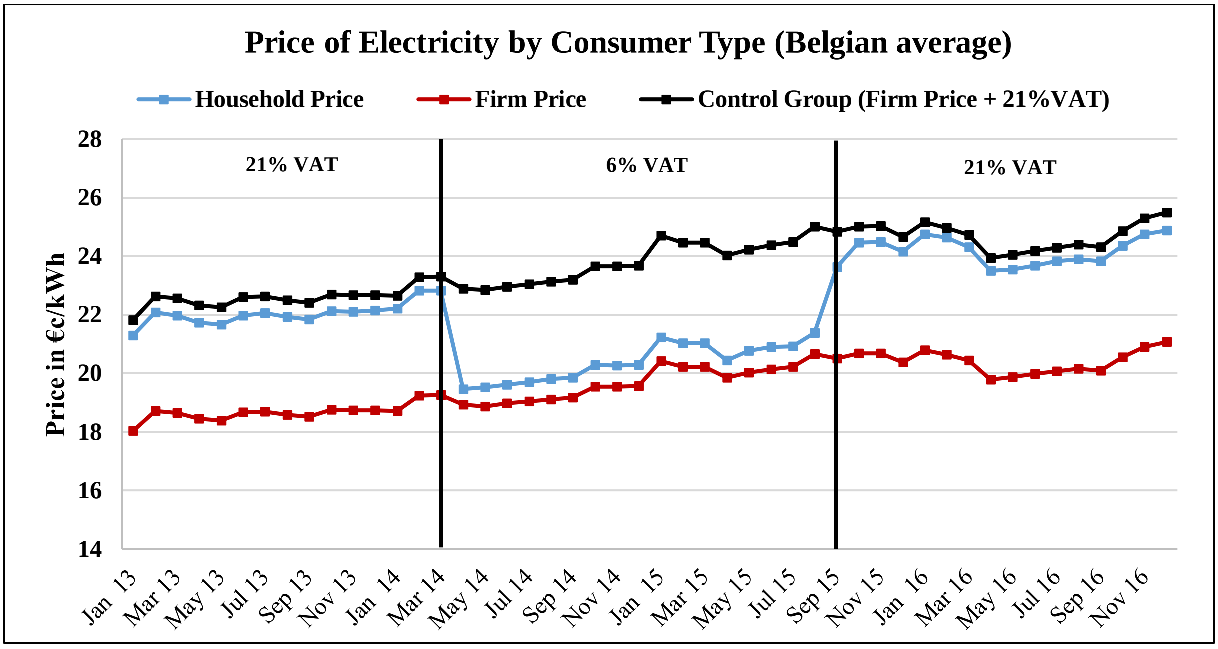

Data source: *Authors with price data from the CREG.*

The period of analysis for the VAT pass-through goes from June 2013 to March 2016. The VAT cut from 21% to 6% occurs in April 2014, and it was then repealed in September 2015. The use of the electricity price for firms as the counterfactual scenario relies on the fundamental assumption that nothing else apart from the temporary VAT cut affected the price of electricity differently across these two groups. Given that the period of the analysis is quite narrow, it is easy to check that no other major policy changes can have affected the price of electricity in Belgium.

Another critical assumption is that the VAT cut only affected the household price. Hence, it did not have an impact on the price of electricity paid by firms (no spillover effect). This assumption seems quite realistic as these two types of electricity services are not substitutable, given that households cannot subscribe to a business tariff. An increase in the household price cannot lead households to switch to a business contract, thereby increasing the electricity price for firms due to higher demand. Furthermore, Figure 4.3 suggests that the electricity price for firms did not vary significantly around the period of policy change. In Figure 4.3, prices are average among service providers. For the tax pass-through analysis, we will use both the cross-sectional and time variations in prices.

Figure 4.4 below shows the evolution of the average electricity consumption of Belgian households during the period 2013-2016. From this graph, we can see that electricity consumption is highly seasonal, with substantial increases in use during the winter periods. The series in black shows the monthly evolution of electricity consumption adjusted for the seasonality in the data. The challenge of estimating the impact of the VAT reform on electricity consumption is to isolate the effect of the tax change from other events that might have occurred over the same period. We do that by running an event study regression controlling for seasonality and other determinants of electricity consumption. We do not perform a difference-in-differences as for the price analysis since we do not have a reliable control group for consumption. Indeed, business electricity consumption is much higher and follows different trends than household consumption. Figure 4.4 reveals average consumption across the 15 service providers. To estimate demand elasticity, we will use both the time and cross-sectional variation in prices and consumptions.

Figure 4.4

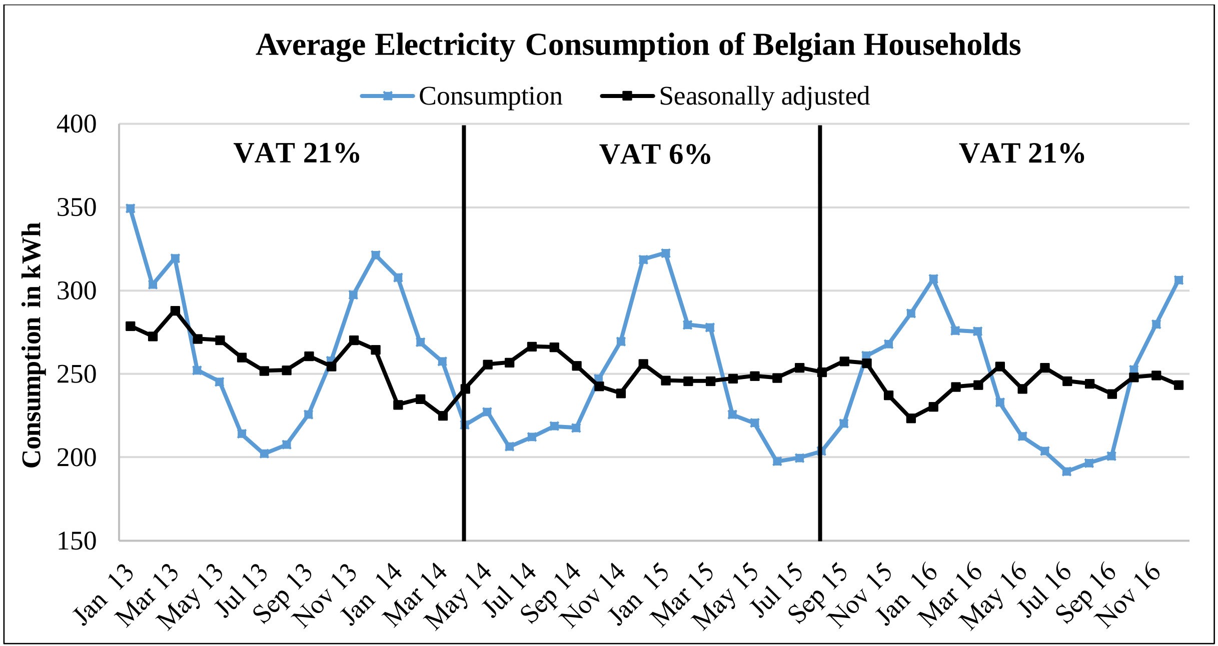

Data source: *Authors, with data from Synergrid. Consumption in kWh is the average across households. Seasonality is computed as the monthly deviation from the average consumption in Belgium over the period.*

The analysis is structured as follows. In the first part, we estimate the VAT pass-through. This analysis is done on both the total electricity price and each price component of the electricity bill. As transmission and distribution have regulated prices, we estimate the pass-through for the price components that are in competitive markets (i.e., the service provider downstream and the electricity producer upstream). We then also test for possible asymmetries in the tax shifting. In the second part, we study the impact of the VAT reform on electricity consumption. The fact that this policy change was exogenous to the consumer demand provides us with the opportunity to estimate the price elasticity of electricity demand. Since consumers experienced a substantial decrease in price and then a similar price increase, we also test for some possible asymmetries in the demand response to these opposite price variations.

4.6.1 VAT Pass-Through to Total Electricity Price

We estimate the impact of the temporary VAT cut on the total price of electricity in Belgium by means of a *difference-in-differences* regression. We use the total price of electricity paid by professional consumers as a control group during the period of the temporary tax change. The model is specified as follows:

$$LN(P)_{ijt} = \beta_0 + \beta_1 HH_j + \beta_2 VAT_{cut} + \beta_3 LN(1+\tau)_t + \varepsilon_{ijt}. \tag{4.1}$$

Where $LN(P)_{ijt}$ is the natural log of the (VAT inclusive) electricity price of provider i for consumer profile j in month t. The consumer profiles j are either households or enterprises. β_0 is the intercept of the model and measures the average (VAT inclusive) price for firms (control group). The variable HH_j is a dummy equal to 1 if the price concerns a household tariff and 0 otherwise. Its coefficient β_1 captures the average difference between the firm price and the household price. VAT_{cut} is a dummy variable equal to 1 during the period of the temporary VAT cut and 0 otherwise. That is, from April 2014 until August 2015. Its coefficient β_2 measures the evolution of the electricity price for firms during this period, which is the counterfactual scenario for household prices. As the dependent variable is expressed in logarithmic terms, this coefficient measures the percentage change in the electricity price for firms during the period of the VAT reduction.

$LN(1+\tau)_t$ is the natural log of the VAT multiplicator, where τ is the current VAT rate in month t. (i.e. either 0.06 or 0.21) The coefficient β_3 can be interpreted as the *difference-in-differences* estimator of the VAT pass-through rate. This coefficient measures the pass-through rate of the temporary VAT reform by considering the evolution of the electricity price for firms as the counterfactual scenario. More specifically, the VAT pass-through rate can be defined as follows:

$$\beta_3 = \frac{\partial LN(P)}{\partial LN(1+\tau)} = \frac{\Delta P/P}{\Delta\tau/(1+\tau)}$$

The model is estimated for the period going from June 2013 until March 2016. Table 4.3 below shows the OLS estimates for the parameters of the *difference-in-differences* regression outlined in equation (1). From these estimates, we can infer the causal impact of the temporary VAT change on electricity prices. The coefficient β_1 shows that household prices are slightly lower than firm prices, including VAT (control group). The coefficient β_2 measures the percentage change in the price of the control group during the period of the temporary VAT cut. Its estimate indicates that firm prices (including VAT) increased by 1% during this period.[67] This means that in case the VAT cut was not implemented, the household prices would have increased by the same amount. The coefficient β_3 indicates the pass-through rate of the temporary VAT reform. The estimation suggests that the temporary VAT change was entirely shifted to electricity prices, with a pass-through rate equal to 100%.

Table 4.3

Impact of VAT Reform on Total Electricity Price

Variable	**Simple Difference**	**Diff-in-Diffs**
Intercept ($\boldsymbol{\beta_0}$)	2.95*** (0.01)	2.96*** (0.01)
Household Tariff ($\boldsymbol{\beta_1}$)		-0.02*** (0.04)
VAT Cut ($\boldsymbol{\beta_2}$)		0.01*** (0.08)
VAT Reform ($\boldsymbol{\beta_3}$)	0.93*** (0.03)	1.00*** (0.04)
Pass-through Rate	**93%**	**100%**
C.I.	[0.88 – 0.98]	[0.93 – 1.08]
$\boldsymbol{R^2}$	0.27	0.28

Notes: **, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors are in parenthesis. There are 6,256 observations in both model specifications.*

[67] As we specified the model in log terms, we do not find much difference if we adjust or not the control group with the VAT rate. Yet, this delivers different results when estimating tax shifting in absolute terms.

We also re-run regression (4.1) to check whether the perfect shifting conceals some heterogeneity in tax shifting across different service providers and contract types. However, we do not find any tax shifting heterogeneity across providers and contracts. Hence, confirming the result of perfect tax shifting of the VAT change on electricity prices. Table 4.3 also shows the result of a simple difference estimator for the VAT pass-through. As we can see from the results, controlling for the change in business price affects the results. With a simple difference estimator, we would have concluded wrongly that the VAT was slightly (but significantly) under shifted. Although Table 4.3 suggests a parallel trend before the tax changes, we also check this more formally by running two placebo tests where the time of the reforms is changed. The results of these tests are shown in Table 4.A.1 in the appendix, both suggest no change in residential electricity prices when the tax cut and tax hike are set differently.

Table 4.4

Impact of VAT Reform on Electricity Price Components

Variable	**Energy**	**Green Quota**
Intercept ($\boldsymbol{\beta_0}$)	1.82*** (0.06)	0.00 (0.20)
HH Tariff ($\boldsymbol{\beta_1}$)	-0.06** (0.03)	0.00 (0.08)
$\mathbf{VAT_{cut}}$ ($\boldsymbol{\beta_2}$)	0.02 (0.01)	0.02 (0.08)
$\mathbf{VAT_{pass-through}}$ ($\boldsymbol{\beta_3}$)	**1.02***** (0.25)	**1.09** (0.83)
$\boldsymbol{R^2}$	0.24	0.01

Notes: **, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors, clustered at the distributor level, are in parenthesis. The first model specification 'Energy' has 272 observations, while the second 'Green Quota' has 817 observations. This is because the energy cost is homogeneous across regions, while green quota can vary regionally.*

The total electricity price is the sum of different price components. Two of these components have regulated prices as they are provided by network operators that are local monopolies. These are the prices for the distribution and transmission of electricity. Given that the energy regulator regulates these prices, network operators cannot respond freely to the VAT change. As a result, the VAT cut must have been entirely shifted. Service providers, however, are in a competitive market environment. They can choose how to shift the tax change on the price of the electricity that they provide. In particular, service providers charge to consumers the cost of the energy supplied and the *green quota* for renewable energy. The VAT may have been shifted differently on these two price components. Therefore, we re-run equation (1) to estimate the VAT pass-through for these two components separately. The results are displayed in Table 4.4 above. Service providers have fully shifted the VAT to the cost of energy with a pass-through rate of 102%. The impact of the VAT on the *green quota* is instead noisier, as the share of this component on the electricity bill is the smallest in magnitude. We estimate an average pass-through of 109%, which is not statistically significant. Therefore, not regulated suppliers in the electricity market have reacted similarly to those that are regulated (i.e., the DSOs).

4.6.2 Symmetry in VAT Pass-through

The temporary VAT reform provides an opportunity to test for asymmetries between the shifting of the tax cut and the consecutive tax hike. Although the results above suggest a full shifting of the VAT reform, we check whether this finding still holds when studying the two events separately. That is, we would like to test whether the VAT cut in April 2014 and the successive VAT hike in September 2015 were both perfectly shifted to electricity prices. To test for this hypothesis, we estimate a *difference-in-differences* regression when the two VAT reforms are considered separately. The model estimated is as follows:

$$LN(P)_{ijt} = \beta_0 + \beta_1 HH_j + \beta_2 VAT_{cut} + \beta_3 LN(1+\tau)_t + \\ + \beta_4 VAT_{hike} + \beta_5 (VAT_{hike} * HH_j) + \varepsilon_{ijt}. \qquad (4.2)$$

Regression (4.2) is similar to regression (4.1). $LN(P)_{ijt}$ is the natural log of electricity price of provider i for consumer profile j in month t. β_0 is the average price in the control group during the period before the temporary VAT reform. While β_1 measures the difference between this price and the average household price during the same period. The interpretation of β_2 and β_3 are the same as in regression (4.2), with the only difference that the baseline period here is only the one before the VAT cut (June 2013 – March 2014).[68] The VAT_{hike} is a dummy variable equal to 1 for the period after the VAT hike. Its coefficient β_4 indicate the percentage price difference in the control group from the period before the VAT reform and the period after the VAT reinstatement at 21%. This would be the counterfactual scenario for household prices if the VAT change did not occur.

The coefficient β_5 is the *difference-in-differences* estimator of the VAT reinstatement at 21% on household electricity prices. Given that we are using the time before the temporary VAT

[68] In regression (4.2) the baseline period was the period outside the VAT cut: before April 2014 and after September 2015.

reform as the baseline period, β_5 is testing whether household electricity prices revert to its initial before the reform. If that is the case, the coefficient β_5 should be equal to zero as this means that after the VAT reinstatement at 21%, household electricity prices reverted to their initial level (controlling for the evolution of firms prices). The results displayed in Table 4.5 suggest that the impact of the two VAT changes on electricity prices was symmetric and that both tax changes were entirely shifted. The pass-through of the VAT cut on the total electricity price is equal to 99%, while that of the energy components is 102%. Again, the pass-through for the green quota is 107%, but it is not statistically significant.

Table 4.5

Test for Asymmetries in VAT Pass-Through

Variable	Total	Energy	Green Quota
$\mathbf{VAT_{cut}}$ $(\boldsymbol{\beta_2})$	0.05*** (0.00)	0.00 (0.03)	0.14 (0.09)
$\mathbf{VAT_{pass-through}}$ $(\boldsymbol{\beta_3})$	**0.99***** (0.04)	**1.02***** (0.12)	**1.07** (0.96)
$\mathbf{VAT_{hike}}$ $(\boldsymbol{\beta_4})$	0.04*** (0.00)	-0.10*** (0.04)	0.14 (0.10)
HH Tariff * $\mathbf{VAT_{hike}}$ $(\boldsymbol{\beta_5})$	0.00 (0.01)	-0.00 (0.06)	0.01 (0.16)
$\boldsymbol{R^2}$	0.34	0.29	0.02

Notes: **, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard error, clustered at the distributor level, are in parenthesis.*

4.6.3 The Impact of the VAT Reform on Electricity Demand

We study the impact of the VAT reform on electricity consumption by measuring the change in electricity demand during the policy change while controlling for other confounding factors. The period of study is four years, from January 2013 to December 2016. During this period, we first observe a VAT decrease from 21% to 6% in April 2014 and then a VAT reinstatement at 21% in September 2015. Therefore, the aim is to estimate the percentage change in electricity demand following these two events. More specifically, we run the following regression.

$$LN(C)_{jt} = D_j + M_t + \beta_C VAT_{t>cut} + \beta_H VAT_{t>hike} + \sum_n \theta_n X_{njt} + \varepsilon_{jt}. \tag{4.3}$$

Where $LN(C)_{jt}$ is the natural log of the (average) electricity consumption (measured in kWh per capita) in the area supplied by distributor j during month t. D_j is the distributor fixed effect, which measures the average consumption level in the area supplied by distributor j. M_t is the month fixed effect that captures seasonality at the monthly level. They correspond to one dummy variable for each month of the year (11 dummies with January being the reference month). The variable $VAT_{t>cut}$ is a dummy variable equal to 1 if the monthly observation is during or after the first VAT change occurred in April 2014. Its coefficient β_C measures the percentage change in electricity demand after the VAT cut. The variable $VAT_{t>hike}$ is a dummy variable equal to 1 if the monthly observation is during or after the VAT hike in September 2015. Its coefficient β_H captures the percentage change in electricity demand after the VAT hike. The baseline period is the one before the first VAT reform. Hence, if the two effects cancel out (i.e., symmetric response), we would observe that $\beta_C + \beta_H = 0$. Lastly, X_n is the set of control variables included in the regression that can vary both at the distributor and monthly level. These variables include the amount of electricity auto produced by households, the degree-day (*Degrés-jours Synergrid*), and the hours of sunlight during that month. We also include the electricity price net of VAT, so that we can control for other price changes not related to the VAT reform.

We run this regression for both standard and social tariffs. Social tariffs are homogenous across the country and are entirely regulated.[69] They can be changed every quarter. Yet, they are also subject to VAT, and hence they have been impacted by the temporary VAT reform (VAT change was entirely shifted on social tariffs). These tariffs are entitled to low-income households under a specific set of legal conditions. As we can identify the electricity consumption subject to these tariffs, we can estimate the demand response of poorer households, and compare it to the rest of the population. The results of regression (4.3) are shown in Table 4.6 below. Regression (4.3) measures the change in demand relative to the baseline demand before the VAT change (January 2013 – March 2014). Columns (1) & (2) and (5) & (6) of Table 4.6 also show the result of a different version of regression (4.3), where we measure the change in electricity demand by using as baseline period the one before (January 2013 – March 2014) and after (September 2015 – December 2016) the VAT cut, where the VAT was at 21%.

From column (1) to column (4) of Table 4.6, we show the estimated coefficients for households subject to standard tariffs, which account for the majority of electricity consumption. The first two columns show that during the period of the VAT cut at 6%, electricity demand was around

[69] Although these tariffs are homogeneous across regions, the final price paid by consumers subject to social tariffs can vary due to some additional regional surcharges. Like the *Cotisation Fonds Energie* in Flanders and the *redevance de raccordement* in Wallonia.

2% larger than the period when the VAT as at 21%. This suggests that the VAT reduction has increased electricity demand. Columns 3 and 4 display a similar result by separating the demand response to the VAT cut and the successive VAT hike. The demand response to the two VAT changes looks very symmetric. Households first increased their demand by 2% when the VAT was reduced at 6% and then decreased their demand to the same extent when the VAT was reinstated at 21%. As a result, after the VAT reinstatement, their demand reverts to the level before the VAT cut (controlling for climate variables and auto production level). Thus, for the same price change, but of a different sign, we get a symmetric change in demand.

Table 4.6

Impact of VAT reform on demand

	Standard Tariffs (1-4)				**Social Tariffs (5-8)**			
Variable	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
VAT cut (β_C)	**2.12** (0.20)	**2.05** (0.20)	**2.06** (0.41)	**2.07** (0.40)	**1.99** (0.22)	**1.81** (0.23)	**2.59** (0.49)	**2.48** (0.49)
VAT hike (β_H)			**-2.14** (0.23)	**-2.05** (0.24)			**-2.08** (0.27)	**-1.60** (0.28)
Degree-Day	0.56 (0.02)	0.57 (0.02)	0.56 (0.02)	0.57 (0.02)	0.56 (0.02)	0.59 (0.02)	0.54 (0.02)	0.58 (0.02)
Sunlight	-0.49 (0.03)	-0.48 (0.03)	-0.49 (0.03)	-0.48 (0.03)	-0.47 (0.04)	-0.44 (0.04)	-0.44 (0.04)	-0.44 (0.04)
Auto production	-1.59 (0.05)	-1.50 (0.07)	-1.57 (0.12)	-1.51 (0.12)	-0.77 (0.07)	-0.72 (0.07)	-0.98 (0.16)	-0.93 (0.16)
Price (no VAT)		-0.04 (0.02)		-0.04 (0.02)		0.34 (0.05)		0.35 (0.05)

Notes: *Standard errors, clustered at the distributor level, are in parenthesis. All coefficients are statistically significant at the 0.01 level, except the one on price (net of VAT) in columns (2) and (4). All models have 720 observations and include interactions between the distributor and monthly FEs (seasonal effect at the distributor level). Their R-squared is around 99%. Without including such interactions, results are very similar, but the R-squared drop to 97%.*[70] *The dummy variables VAT cut, and VAT hike are rescaled (divided by 100) so that they can present the coefficients in percentage points.*

Columns (5) to (8) of Table 4.6, display the estimated coefficients for households subject to social tariffs. They account for a minority of the electricity demand, which is composed of low-income households mostly. Results are similar to those of other households, with their change

[70] The R-squared of the model is very high, which can suggest some overfitting issues. Nevertheless, this is because electricity consumption is highly seasonal in every area. In fact, once removing seasonal dummies and distributor fixed effects, the R-squared drops from 99% to roughly 6%.

in electricity demand close to 2%. Yet, the last two columns show that demand increased by more than 2% during the VAT cut, while it dropped by less than 2% after the VAT reinstatement. However, after January 2016, social tariffs were reduced by the energy authority (partly to protect the low-income group from the VAT reinstatement). This social tariff adjustment can explain why demand did not revert to the level before the VAT cut. Nevertheless, including the price net of VAT as a control variable, does not account for this effect (its coefficient is even positive).[71]

The coefficients for the control variables included in regression (4.3) have all the expected signs. The degree-day measures how much heating is needed over a given month as a function of the outside temperature. The sign of its coefficient is positive and significant, which means that electricity consumption increases when more heating is needed. The coefficient for the monthly hours of sunlight is instead negative and significant. This indicates that more hours of sunlight lead to lower electricity consumption, which is possibly due to reduced home lighting. Lastly, the auto production of electricity tends to reduce electricity consumption from the grid. This result is also quite intuitive as these two sources of energy are substitutes. Most of the auto production of electricity in Belgium is through solar panels, which allows households to consume their own energy instead of getting it from the electricity grid.

4.6.4 The Price Elasticity of Electricity Demand

The exogenous variation in electricity prices due to the VAT reform allows us to estimate the price elasticity of electricity demand. The estimates can be easily retrieved from regression (4.3), as we know the change of both price and demand after the VAT reform. However, the VAT reform was the principal but not the only source of price variation in the data. Changes in electricity prices over the sample period also occur due to updates in regulated components, such as transmission and distributions, and due to some other shocks to energy prices. As we can disentangle the type of price variations, we can estimate the demand elasticity by exploiting different sources of changes in price.

The change in price due to the VAT change is very likely to be exogenous to household demand because both reforms were implemented for reasons outside the electricity market. This constitutes a good source of price variation to identify the demand elasticity. Another source of

[71] Social tariffs are fully regulated and the only two sources of variation over the sample period are the VAT change (mainly) and the update in January 2016. Thus, the positive coefficient for the price net of VAT should not be indicative, as it could be due to the little variation in the price data and its correlation with some unobserved factor after January 2016.

price variation is the change in regulated price components. The energy regulator sets prices for transmission and distribution of electricity, which change at least once per year. Over the sample period, they steadily increased. These prices should be cost-based and not related to changes in demand-side conditions. For this reason, they can also constitute a useful source of price variation to identify the demand elasticity. The residual variation in electricity prices is due to changes in the cost of energy, which is not regulated. These price changes can be due to temporal or permanent cost shocks in producing and selling electricity to residential consumers. Nevertheless, they can also be related to changes in demand-side conditions, as service providers are free to set their prices. In such a case, price changes are not fully exogenous to household demand and can, therefore, constitute a spurious source of price variation.

We estimate the demand elasticity by exploiting these different sources of price variation. In particular, we identify the demand elasticity by using either: i) the VAT change; ii) the VAT change plus the regulated price change; or iii) all price changes. We do this exercise for both standard and social tariffs.[72] More specifically, we estimate the demand elasticity by running the following regression:

$$LN(C)_{jt} = D_j + M_t + \beta LN(Price)_t + \sum_n \theta_n X_{njt} + \varepsilon_{jt}. \tag{4.4}$$

This regression is mainly similar to regression (4.3) except that the VAT dummies are replaced directly by the price variable $LN(Price)_t$. The price variable is either: i) the log of the VAT rate, $LN(1+\tau)_t$, when we exploit the VAT change only; ii) the log of the regulated price (VAT included); or iii) the total electricity price. The coefficient β is the price elasticity of electricity demand. We estimate regression (4.4) by standard OLS procedure. Moreover, we also estimate the first two specifications through a two-stage least square (2SLS) procedure. Where in the first stage, we instrument the price by either a dummy variable for the VAT cut or by the regulated price (net of VAT) plus the VAT dummy. The results of this analysis are displayed in Table 4.7 below. The β coefficient for price indicates the estimate for the demand elasticity. According to the different model specifications displayed in Table 4.7, this can range from -0.09 to -0.17. Although these estimates seem quite low, they are nevertheless significant and in line with other studies based on quasi-experimental price variations (see Ito, 2014; Deryugina, MacKay, & Reif, 2020). Interestingly, Table 4.7 highlights that (poor) households subject to

[72] However, as social tariffs are fully regulated, estimating the demand elasticity using the VAT change and the regulated price component is equivalent to exploiting all price variation in the data.

social tariffs have a similar demand elasticity than other consumers. When exploiting all price variation, this is equal to -0.14 and -0.12, respectively (see columns 5 and 8).

Table 4.7

Demand Elasticity

	Standard tariffs (1-5)					**Social tariffs (6-8)**		
	VAT only		**VAT and regulated price**		**Full price**	**VAT only**		**Full price**
Variable	**OLS (1)**	**2SLS (2)**	**OLS (3)**	**2SLS (4)**	**OLS (5)**	**OLS (6)**	**2SLS (7)**	**OLS (8)**
LN(price)	**-0.16** (0.01)	**-0.17** (0.03)	**-0.09** (0.01)	**-0.10** (0.01)	**-0.12** (0.02)	**-0.17** (0.02)	**-0.17** (0.02)	**-0.14** (0.02)
Degree-Day	0.56 (0.02)	0.50 (0.04)	0.56 (0.02)	0.52 (0.03)	0.58 (0.02)	0.55 (0.02)	0.55 (0.02)	0.53 (0.02)
Sunlight	-0.48 (0.03)	-0.33 (0.05)	-0.47 (0.03)	-0.30 (0.03)	-0.46 (0.03)	-0.45 (0.04)	-0.44 (0.04)	-0.45 (0.04)
Auto production	-1.59 (0.05)	-1.43 (0.08)	-1.21 (0.05)	-1.14 (0.05)	-1.32 (0.05)	-0.85 (0.07)	-0.85 (0.07)	-0.86 (0.07)
Elasticity	**-0.16**	**-0.17**	**-0.09**	**-0.10**	**-0.12**	**-0.17**	**-0.17**	**-0.14**
C.I.	0.13-0.19	0.10-0.24	0.08-0.10	0.07-0.13	0.10-0.15	0.13-0.20	0.14-0.21	0.10-0.17

Notes: *all coefficients are statistically significant at the 0.01 level. Standard errors, clustered at the distributor level, are in parenthesis). All models have 720 observations and include interactions between the distributor and monthly FEs (seasonal effect at the distributor level). Their R-squared is around 99%. Without including such interactions, results are very similar, but the R-squared drops to 97%.*

If we measure the demand elasticity using the VAT change only, we get a slightly larger demand elasticity, for both standard and social tariffs, which are equal to -0.17. Interestingly, when exploiting changes in both the VAT and regulated price components, we get significantly lower demand elasticities. For standard tariffs, this decreases the demand elasticity from -0.17 to -0.10. Using OLS or 2SLS does not really affect the coefficients, but it slightly increases the range of the confidence intervals. Hence, these findings suggest that consumers are more sensitive to changes in the VAT rate than changes in other price components. Two elements could help explain such a result. First, the VAT change was announced and hotly debated in the media, which makes the price variation quite salient to consumers. Second, the change in price was easy to understand, as it mostly consisted of the difference in the VAT rate. In contrast, changes in other price components tend to be smaller and less noticeable, as they have to be carefully checked in the monthly electricity bill. [73]

4.6.5 Symmetry in Demand Elasticities

These policy changes provide us with an excellent natural experiment to test for differences in demand response to a price hike and price cut. The answer to such a question has quite relevant policy implications, as most climate change policies would imply some energy price hike. If consumers have asymmetric responses to price changes of different signs, then this must be taken into account when evaluating the impact of these policies on energy use. Moreover, evidence of the effects of exogenous price hikes on electricity demand is very limited. Table 4.6 has shown that households responded to the VAT cut and the VAT reinstatement symmetrically. This means that price changes of different signs have a similar but opposite effect on demand. Therefore, estimates of demand elasticities obtained by exploiting exogenous decreases in electricity price could also be indicative of the consumer response to a price hike.

More formally, we measure demand elasticity by allowing it to vary across two periods: one that goes from January 2013 to December 2014; another that goes from January 2015 to December 2016. The aim is to estimate demand elasticities for the VAT cut and the subsequent VAT hike separately. Since the VAT reform was not the only source of price variation, we do not have a perfectly symmetric price variation. In particular, there is a price drop of -11% in the first period and a price hike of +14% for standard tariffs and +12% for social tariffs, in the second period. Therefore, we re-estimate regression (4.4) using the full price variation to control for the asymmetry in price changes and compute the demand elasticity separately for the price cut and the price hike. The results are displayed in Table 4.8 below.

Table 4.8

Symmetry in Demand Elasticities

Variable	Standard Tariffs	Social Tariffs
Demand Elasticity (VAT cut)	**-0.10**	**0.12**
C.I.	[0.08 – 0.13]	[0.08 – 0.15]
Demand Elasticity (VAT hike)	**-0.10**	**-0.12**
C.I.	[0.08 – 0.13]	[0.08 – 0.16]

Notes: *all coefficients are statistically significant at the 0.01 level. Standard errors, clustered at the distributor level, are in parenthesis. Control variables include distributor FE, monthly FE, and their interaction, as well as all other control variables included in model (4).*

[73] For instance, Sexton (2015) shows that consumers enrolling in automatic bill payment programs become less price sensitive, increasing their electricity demand due to a decline in price salience.

As already suggested in Table 4.6, consumers are equally sensitive to price cuts and price hikes. The demand elasticity is symmetric. However, the estimates are slightly lower than in Table 4.7. Again, we do not find much difference in demand elasticities across different household types. Most households have a demand elasticity of -0.10, while for (poor) households subject to social tariffs, the elasticity is -0.12. This result can have important implications for climate change policies, as it highlights that consumers are not only sensitive to reductions in electricity prices. They are also equally sensitive to price increases due to taxation.

4.6.6 Regional and Time Heterogeneity

We further investigate the extent of heterogeneity in demand elasticity across Belgian areas. We re-estimate regression (4.4) by allowing each Belgian region to have different demand elasticities. The results of this estimation lead us to compute the demand elasticities for the Region of Brussels, Flanders, and Wallonia. These estimates are shown in Table 4.9, where we use the full price variation.[74] Concerning standard tariffs, households in Wallonia and Flanders display similar demand elasticities in the range of -0.13 and -0.11. Households in the Region of Brussels display a higher demand elasticity of -0.20, except for those subject to social tariffs who display the lowest elasticity, not significantly different from zero. Households entitled to social tariffs in Flanders and Wallonia have a demand elasticity of -0.20 and -0.10, respectively. This evidence thus suggests that similar price changes can have a different impact on electricity demand across geographical areas and household types. Residents in the urban area of Brussels and Flemish households subject to social tariffs are the most responsive to changes in prices and hence energy taxes. In contrast, Brussels residents entitled to social tariffs do not respond to price changes significantly.

Table 4.9

Regional Heterogeneity in Demand Elasticity

	Region of Brussels	**Flanders**	**Wallonia**
Standard Tariffs	**-0.20**	**-0.13**	**-0.11**
C.I.	[0.16 – 0.24]	[0.08 – 0.18]	[0.03 – 0.18]
Social Tariffs	**-0.03**	**-0.20**	**-0.10**
C.I.	[-0.13 – 0.07]	[0.15 – 0.24]	[0.04 – 0.15]

Notes: *Demand elasticities are computed at the sample average. Estimates are obtained through the estimation of regression (4.4) by allowing for regional heterogeneity in demand response.*

[74] The estimates are very similar when computing elasticities using only VAT or the VAT plus regulated price changes.

We further test for the time heterogeneity in the demand response to changes in electricity price. According to the literature, the demand response to changes in electricity prices is dynamic (Labandeira, Labeaga, & López-Otero, 2017; Deryugina, MacKay, & Reif, 2020). The demand elasticity tends to increase over time as consumers may take time to change their stock of appliances and their relative energy efficiency (Rapson, 2014; Sahari, 2019). Another possible explanation is that electricity prices are complex and not easily visible to consumers (Sexton, 2015). Thus, consumers may take time to realize an actual price change. In our context, the price change is likely to be more salient than a usual price change for at least two reasons. First, the VAT reforms were publicly announced in the media. Second, it was easy to assess the amount of the price change, given that it is close to the variation in the VAT rates. This means that the demand response to the VAT reform could have been faster than other price changes.

To explore the dynamics in the demand response, we estimate the variation in the kWh of electricity consumed during each month of VAT at 6% by taking as reference period the one with the VAT at 21%. In practice, we estimate the following regression.

$$C_{jt} = D_j + M_t + \sum_{t=-6}^{16} \beta_t (VAT_{cut})_t + \sum_n \theta_n X_{njt} + \varepsilon_{jt}. \tag{4.5}$$

Where C_{jt} is the electricity consumption measured in kWh per capita for the area supplied by distributor j during month t. We do not specify the dependent variable in log terms as we want to estimate the absolute change in kWh consumed during the period of the reform. VAT_{cut} is a dummy variable equal to 1 for each month before and after the implementation of the VAT cut. We allow t to be $t \in [-6, 16]$, which means that we include six leads and 16 lags from the month of the VAT cut, which occurs at $t = 0$.[75] The coefficient β_t measures the increase in kWh consumed over each month before and after the VAT reduction at 6%. The reference period is the one with the VAT at 21%, with the exclusion of the six months pre-reform (as we need them to estimate the leads).

By including the leads up to six months before the VAT cut, we can estimate whether electricity demand increased even before the VAT was reduced. This is a useful robustness check as we know that VAT change was only announced one month before its implementation. Thus, we should not observe any effect due to the reform in those months. The lags instead indicate how the variation in electricity demand evolved over each month of VAT at 6%. All other variables

[75] The last lag coincides with the month just before the reinstatement of the VAT at 21%.

included in regression (4.5) are the same as in regression (4.3). The results of this analysis are displayed graphically in Figure 4.5 below.

Figure 4.5

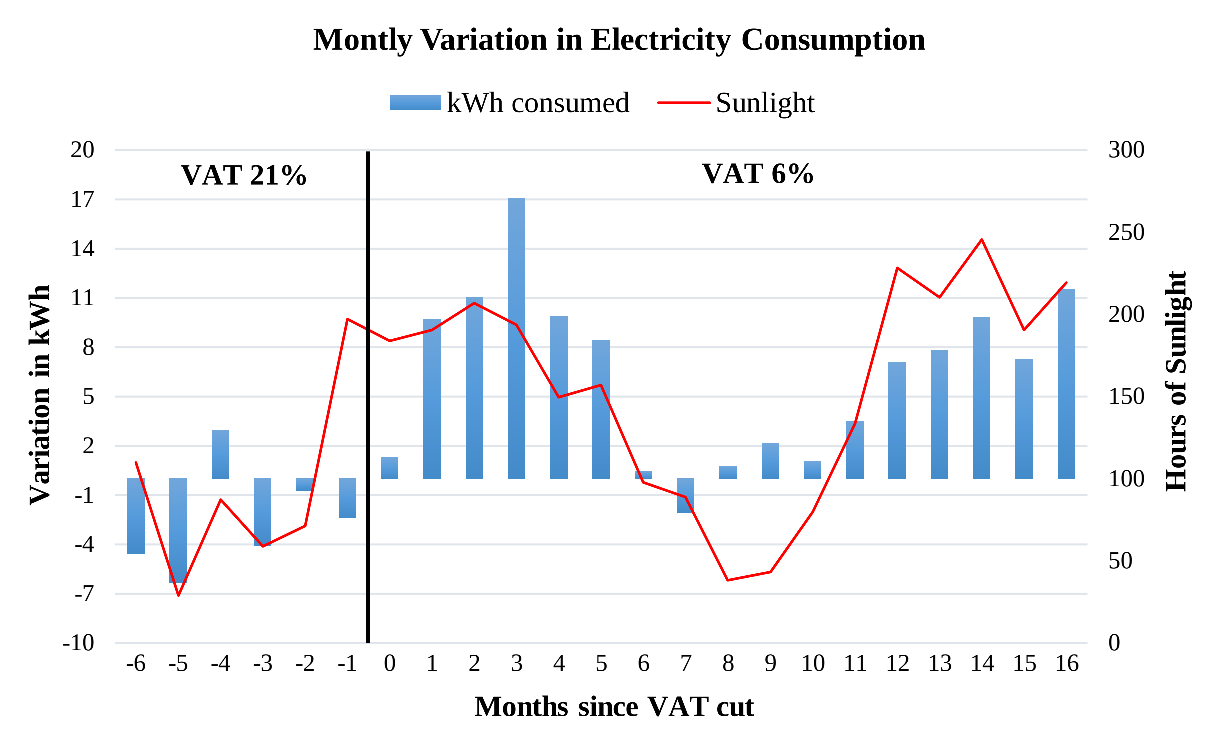

Notes: *the histogram shows the estimates of each monthly* β_t, *which are the leads and lags of regression (4.5). The red curve displays the hours of sunlight during each month.*

Figure 4.5 shows that electricity consumption increased almost immediately after the VAT cut, precisely one month after its implementation. The average monthly increase in electricity consumption during the period of VAT at 6% was equal to 6.30 kWh, which represents roughly a 2% increase in electricity demand. Interestingly, the demand response seems to be frontloaded, with a sharper increase during the first six months of tax reduction. This evidence suggests that consumers were well aware of the price cut and thus increased demand quickly. Subsequently, there is no significant change in electricity demand in the period that goes from October 2014 to February 2015, while the demand increases again in the successive months. Interestingly, Figure 4.5 highlights the strong correlation between the hours of sunlight and the demand response to the VAT change. Most of the demand increase occurs during months that are sunnier (and warmer). A possible reason why the electricity use did not also rise during colder and less sunny months could be because energy consumption was already high during that period. As a result, there was less scope for increasing it at the margin.

4.7 Robustness Checks

For the VAT pass-through analysis, we dispose of a reliable control group to check for the counterfactual evolution of electricity prices. We cannot adopt the same approach to study the demand response. This is because the electricity consumption of SMEs cannot be easily compared to households. SMEs consume much more electricity on average and can be affected by other unobserved factors that are less relevant for households. Nevertheless, to test the robustness of our results, we run a placebo test on the evolution of electricity consumption for SMEs. The idea is to check how the electricity consumption of SMEs varied during the period of the VAT cut. Firms are not subject to VAT and hence should not have been affected by the reform. The test consists of checking whether we observe a similar evolution of electricity demand for SMEs and households, which would then suggest that other factors outside the VAT changes might drive our results. To do so, we re-run regression (4.3) for both household and business consumers. The results of this test are shown in Table 4.10 below.[76]

Table 4.10

Placebo Test: Impact of VAT Reforms on Business Consumption

Variable	Business (1)	Business (2)	Households (3)	Households (4)
VAT_{cut}	**-0.28** (1.33)	**-0.19** (1.11)	**2.07***** (0.40)	**2.01***** (0.42)
VAT_{hike}	**-5.30***** (0.60)	**-1.27** (1.05)	**-2.05***** (0.24)	**1.33***** (0.15)
Auto production	-0.11*** (0.02)	-0.07*** (0.01)	0.57*** (0.02)	-0.18*** (0.01)
Degree-Day	0.04*** (0.00)	0.05*** (0.00)	-0.48*** (0.03)	0.07*** (0.00)
Sunlight	-0.07*** (0.01)	-0.06*** (0.00)	-1.51*** (0.12)	-0.04*** (0.00)
Price (net VAT)	0.04 (0.03)	0.07** (0.03)	-0.04** (0.02)	-0.02 (0.02)
Time Trend	no	yes	no	yes
R-squared	0.98	0.99	0.99	0.99

Notes: **, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level, respectively. Standard errors, clustered at the distributor level, are in parenthesis. Control variables also include distributor FE, monthly FE, and their interaction.*

[76] We also run regression (4.3) by adding additional controls that could determine demand for electricity beside prices. These are disposable income and industrial production. Results of this estimation are in Table 4.A.2 in the appendix. Adding these variables does not change our results.

Column (1) of Table 4.10 shows the results of the placebo test. The coefficient for the VAT cut is equal to 0, indicating that there are no significant changes in electricity demand for SMEs after the VAT cut. During the same period, however, the electricity demand of households increased by 2% (see column 3). Given that only households are subject to the VAT, this evidence suggests that we are capturing the causal impact of the VAT cut on households' consumption. However, the coefficient for the VAT hike is strongly negative and significant. This indicates that the electricity consumption of SMEs has decreased mainly during the period of VAT reinstatement at 21%. This effect is probably due to some other factors affecting SMEs that are unrelated to the VAT change. Indeed, once we account for a time trend in the regression, such effect disappears. These results are shown in column (2), and thus suggests that the electricity consumption of SMEs tended to decline over the period. We also report the estimates for household consumers when including a time trend in electricity consumption. Interestingly, we find that our estimate for the VAT cut is unchanged at 2%. This result suggests that our demand estimates do not capture a pre-existing consumption trend in the data.[77]

4.8 Conclusions

This paper studies the impact of a temporary VAT cut on the Belgian electricity market, which occurred in the period between April 2014 and September 2015. The Belgian government temporarily reduced the VAT rate on electricity prices from 21% to 6%. We study how such VAT cut affected both the price and demand for electricity. We estimate the VAT pass-through to electricity prices by means of a *difference-in-differences* analysis. We use the electricity price paid by firms (not subject to VAT) as a control group for the period before and after the temporary VAT cut. The results show that the VAT cut was fully shifted to electricity prices so that households could entirely benefit from the tax reduction. We do not find any asymmetry in tax shifting between the VAT cut to 6% and the subsequent reinstatement of the VAT to 21%.

We then exploit the exogenous tax reform to estimate the impact of the temporary VAT cut on the electricity demand of Belgian households and get a measure for the demand elasticity. We find that the VAT cut generated a 2.05% increase in electricity demand and that demand elasticity is equal to 0.12, with some heterogeneity across Belgian regions. Furthermore, we provide novel evidence about the symmetry in the consumers' response to price cuts and price hikes. That is, consumers are equally sensitive to a price cut and a price hike of a similar extent. This finding has important implications for price-based climate policies, as it indicates that

[77] The coefficient for the VAT hike is positive instead of being equal to -2%. However, this is probably because the effect of the VAT increase is fully absorbed by the negative time trend in the last periods of our sample.

price hikes can be a useful tool in reducing carbon emission due to electricity generation. We also show that consumers reacted immediately to the VAT cut, increasing their demand the next month. This evidence contrasts with the existing literature, suggesting that the tax change was more salient compared to other price changes.

References

Alberini, A., & Filippini, M. (2011). Response of residential electricity demand to price: The effect of measurement error. *Energy economics*, *33*(5), 889-895.

Benedek, D., De Mooij, R. A., Keen, M., & Wingender, P. (2019). Varieties of VAT pass through. *International Tax and Public Finance*, 1-41.

Berardi, N., Sevestre, P., Tepaut, M., & Vigneron, A. (2016). The impact of a 'soda tax' on prices: evidence from French micro data. *Applied Economics*, *48*(41), 3976-3994.

Besley, T. J., & Rosen, H. S. (1999). Sales Taxes and Prices: An Empirical Analysis. *National Tax Journal*, *52*(2), 157-78.

Boccard, N. & Gautier, A. (2020). *Solar rebound: the unintended consequences of subsidies*. CORE Discussion Papers 2020002, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Bureau Fédéral du Plan. (2015). *Évaluation de la baisse de la TVA sur l'électricité à usage domestique*. Downloaded on 2020 April 12. Available at <https://www.plan.be/uploaded/documents/201507231617290.REP_EVABTWELEK715_11058_00000003.pdf>

Carbonnier, C. (2007). Who pays sales taxes? Evidence from French VAT reforms, 1987–1999. *Journal of Public Economics*, *91*(5-6), 1219-1229.

Cawley, J., & Frisvold, D. E. (2017). The Pass-Through of Taxes on Sugar-Sweetened Beverages to Retail Prices: The Case of Berkeley, California. *Journal of Policy Analysis and Management*, *36*(2), 303-326.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, 99 (4), 1145-77.

DeCicca, P., Kenkel, D., & Liu, F. (2013). Who pays cigarette taxes? The impact of consumer price search. *Review of Economics and Statistics*, *95*(2), 516-529.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature*, *47*(2), 315-72.

Deryugina, T., MacKay, A., & Reif, J. (2020). The long-run dynamics of electricity demand: Evidence from municipal aggregation. *American Economic Journal: Applied Economics*, *12*(1), 86-114.

Doyle Jr, J. J., & Samphantharak, K. (2008). $2.00 Gas! Studying the effects of a gas tax moratorium. *Journal of public economics*, *92*(3-4), 869-884.

Etilé, F., Lecocq, S., & Boizot-Szantai, C. (2019). *The Incidence of Soft-Drink Taxes on Consumer Prices and Welfare: Evidence from the French " Soda Tax".* PSE Working Papers.

Eurostat (2020). Energy (nrg). Energy statistics - quantitis [Data file]. Downloaded on 2020 April 7 (last update: 2020-02-24). Available at <http://ec.europa.eu/eurostat/data/database>.

Fabra, N., & Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, *104*(9), 2872-99.

Gautier, A., & Jacqmin, J. (2020). PV adoption: the role of distribution tariffs under net metering. *Journal of Regulatory Economics*, *57*, 53–73.

Harding, M., Leibtag, E., & Lovenheim, M. F. (2012). The heterogeneous geographic and socioeconomic incidence of cigarette taxes: Evidence from Nielsen Homescan Data. *American Economic Journal: Economic Policy*, *4*(4), 169-198.

Hindriks, J., & Serse, V. (2019). Heterogeneity in the tax pass-through to spirit retail prices: Evidence from Belgium. *Journal of Public Economics*, *176*, 142-160.

Hintermann, B. (2016). Pass-through of CO2 emission costs to hourly electricity prices in Germany. *Journal of the Association of Environmental and Resource Economists*, *3*(4), 857-891.

Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review*, *104*(2), 537-63.

Kenkel, D. S. (2005). Are alcohol tax hikes fully passed through to prices? Evidence from Alaska. *American Economic Review*, *95*(2), 273-277

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

Labandeira, X., Labeaga, J. M., & López-Otero, X. (2017). A meta-analysis on the price elasticity of energy demand. *Energy Policy*, *102*, 549-568.

Marion, J., & Muehlegger, E. (2011). Fuel tax incidence and supply conditions. *Journal of public economics*, *95*(9-10), 1202-1212.

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

Rapson, D. (2014). Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior. *Journal of Environmental Economics and Management*, *68*(1), 141-160.

Sahari, A. (2019). Electricity prices and consumers' long-term technology choices: Evidence from heating investments. *European Economic Review*, *114*, 19-53.

Sexton, S. (2015). Automatic bill payment and salience effects: Evidence from electricity consumption. *Review of Economics and Statistics*, *97*(2), 229-241.

Shrestha, V., & Markowitz, S. (2016). The Pass-Through of Beer Taxes to Prices: Evidence from State and Federal Tax Changes. *Economic Inquiry*, *54*(4), 1946-1962.

4.A Appendices

4.A.1 Additional Figures and Tables

Table 4.A.1

Placebo Test for the VAT Pass-through

Variable	VAT change 10/2013	VAT change 01/2015
Intercept ($\boldsymbol{\beta_0}$)	2.90*** (0.01)	3.58*** (0.00)
Household Tariff ($\boldsymbol{\beta_1}$)	0.17*** (0.01)	-0.59*** (0.04)
VAT Cut ($\boldsymbol{\beta_2}$)	0.02*** (0.01)	0.06*** (0.08)
VAT Reform ($\boldsymbol{\beta_3}$)	-0.01 (0.01)	0.01 (0.01)
Sample Period	06/2013 – 03/2014	05/2014 – 07/2015
$\boldsymbol{R^2}$	0.22	0.24

Notes: **, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors are in parenthesis.*

Table 4.A.2

Robustness Checks: Additional Control Variables

Variable	Standard Tariffs (1)	Standard Tariffs (2)	Social Tariffs (3)	Social Tariffs (4)
$\mathbf{VAT_{cut}}$	**1.82***** (0.49)	**1.83***** (0.49)	**2.89***** (0.69)	**2.96***** (0.68)
$\mathbf{VAT_{hike}}$	**0.73*** (0.41)	**0.74*** (0.41)	**0.11** (0.58)	**0.11***** (0.58)
Auto production	-0.16 (0.20)	-0.14 (0.22)	-0.18 (0.31)	-0.39 (0.31)
Degree-Day	0.06*** (0.00)	0.06*** (0.00)	0.06*** (0.00)	0.06*** (0.00)
Sunlight	-0.04*** (0.00)	-0.04*** (0.00)	-0.05*** (0.00)	-0.04*** (0.00)
Price (net VAT)		-0.02 (0.03)		0.12 (0.36)
Disposable Income	-3.76*** (0.42)	-3.71*** (0.43)	-2.29*** (0.59)	-1.94*** (0.60)
Industrial Production	-0.19 (0.14)	-0.21 (0.14)	-0.13 (0.19)	-0.23 (0.19)

Notes: **, ** and *** indicate statistical significance at the 0.10, 0.05 and 0.01 level respectively. Standard errors are in parenthesis. Disposable income and industrial production are quarterly data at the national level. R-Squared coefficients are similar to those in Table 4.6.*

Figure 4.A.1

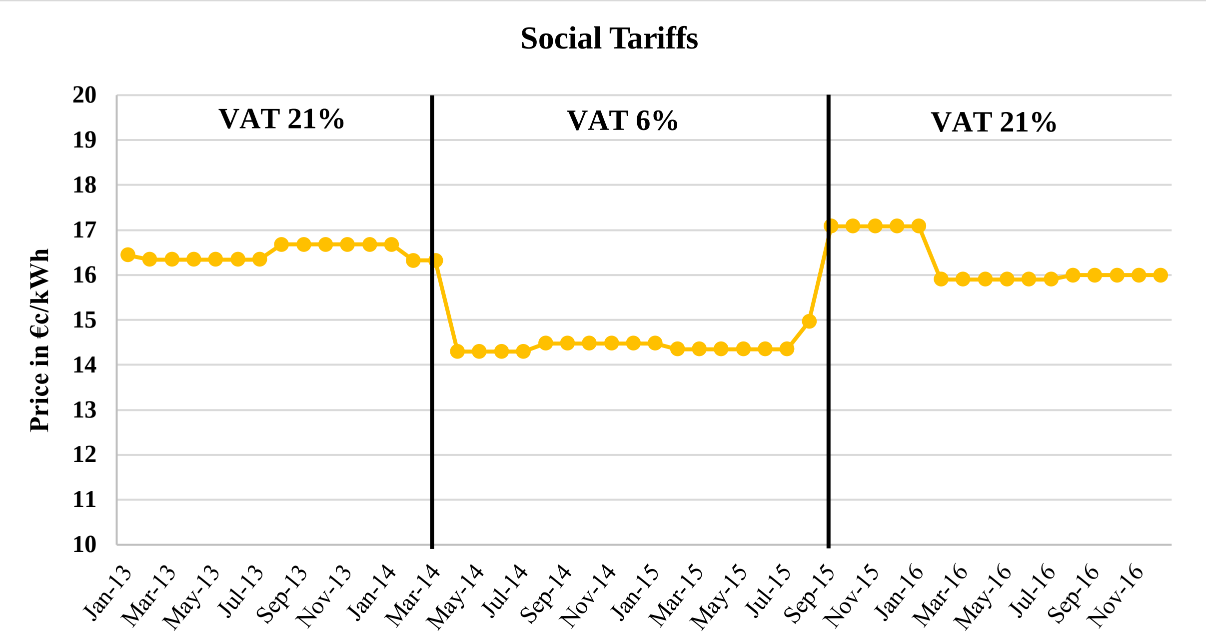

Notes: the series represents the evolution of *social tariffs, which are homogenous across Belgium,*

4.A.2 Lagged Price Model

Although we can only observe a limited time frame, we also measure how demand elasticity varies with time. We re-run regression (4.4) by using monthly lagged prices. We do this for a maximum of 12 months lag and using the full price variation.[78] The results of this analysis are displayed graphically in Figure 4.A.2.

Figure 4.A.2

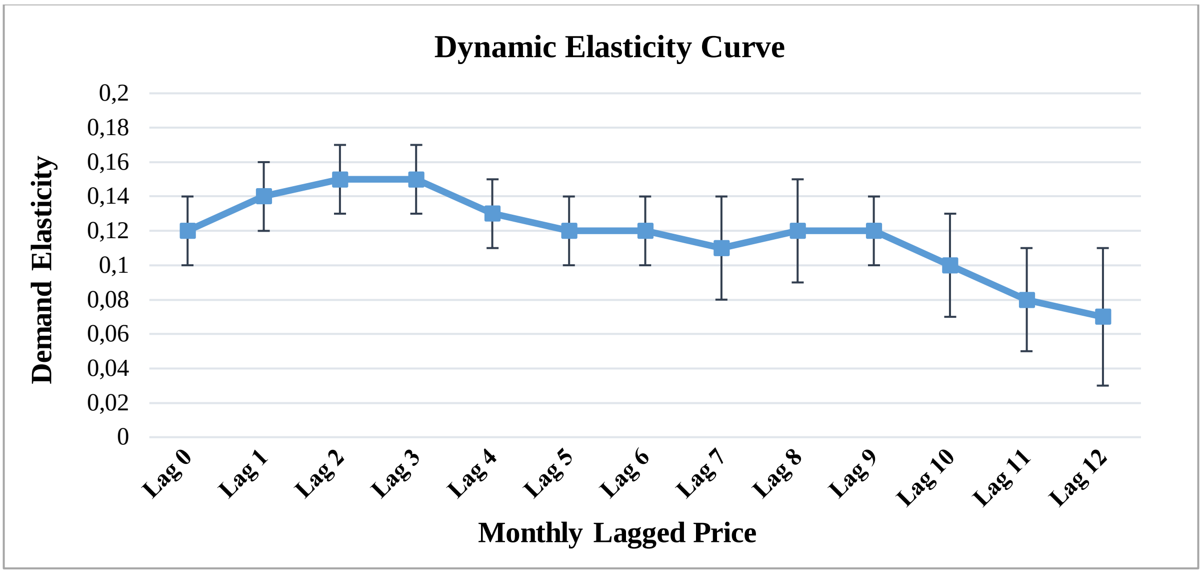

Notes: *the curve shows the demand elasticity computed with different lags and their respective confidence interval. Results are obtained by estimating regression (4.4) with lagged prices instead of the actual monthly price. Each regression for any lag is estimated separately.*

Including lagged prices instead of the actual monthly price does not seem to affect our results significantly. Yet, the larger demand response occurs within two and three months from the price change. This evidence suggests that consumers respond quickly to variations in electricity prices. Furthermore, Figure 4.A.2 highlights a slight negative trend of the demand elasticity over time, which indicates that the demand response to price changes tends to decrease over time. This result is in contrast with existing evidence suggesting that demand elasticity is higher in the long-run. As already discussed above, we are identifying the demand elasticity by exploiting large policy shocks to prices that are more salient than a typical change in electricity price. Hence, this can indicate that consumers respond promptly to electricity price changes if they are quite aware of them.

[78] Results are similar when using different sources of price variation and are available upon request.