

SHOPPING FOR LOWER SALES TAX RATES*

Scott R. Baker[†] Lorenz Kueng[‡]

Preliminary version: please do not cite or circulate

November 2016

Abstract

Using comprehensive high-frequency state and local sales tax data, we find that household spending responds strongly to changes in sales tax rates. Despite their complexity, such as the fact that sales taxes are not observed in posted prices and have a wide range of rates and exemptions, households increase online and cross-border shopping and stock up on storable goods before taxes increase. Interestingly, households adjust spending similarly on both taxable and tax-exempt goods. We demonstrate that this seemingly irrational behavior is optimal in the presence of shopping complementarities and provide independent evidence in favor of this new mechanism. While our results demonstrate that salience of sales tax changes is high on average, we also show that upcoming tax changes that are more salient prompt larger responses.

JEL Classification: D12, H31, E21

Keywords: shopping complementarity, sales tax salience, consumer spending

*We would like to thank Stephanie Johnson for outstanding research and technical assistance. We also thank our discussants Jonathan Parker and Inessa Liskovich, and participants at seminars at the NBER Summer Institute 2016, Arizona State, Berkeley, Bern, Columbia, Frankfurt, Minneapolis Fed, Minnesota, Munich, Northwestern and Yale for their comments. This research received financial support from the Alfred P. Sloan Foundation through the NBER Household Finance small grant program. The results of this paper are calculated based on data from The Nielsen Company (U.S.) LLC and provided by the Marketing Data Center and the University of Chicago Booth School of Business.

[†]Department of Finance, Northwestern University's Kellogg School of Management, 2001 Sheridan Road, Evanston, IL 60208. E-mail: s-baker@kellogg.northwestern.edu.

[‡]Department of Finance, Northwestern University's Kellogg School of Management, 2001 Sheridan Road, Evanston, IL 60208, and National Bureau of Economic Research. E-mail: l-kueng@northwestern.edu.

Most economic analysis rests on the assumption that individuals optimize fully in the presence of limited resources. Increasingly however, research investigates deviations from this paradigm, allowing individuals to make systematic mistakes and to pay limited attention to relevant information or to entire choice dimensions.¹ These deviations could be especially large in the context of public finance where taxes are often complicated and non-transparent. A range of research has found empirical evidence of these deviations from fully rational behavior in responding to a large number of taxes and fees, and recent theoretical work has started to incorporate tax non-salience into models of optimal taxation.²

This paper focuses on sales taxes in the U.S., which are particularly complex and thus provide a powerful setting for detecting potential deviations from full optimization. The complexity of U.S. sales taxation stems from three main sources. First, taxes are typically not included in the posted prices and are only applied at checkout, making it difficult for consumers to take sales taxes into account when choosing different products in a store.³ Second, some goods are exempt from taxation, with different states (and some local jurisdictions) applying different exemption rules. Third, due to strong fiscal federalism in the U.S., a wide range of overlapping tax jurisdictions can impose their own sales taxes (state, county, city, and special districts), leading to a large range of tax rates across geographic locations down to the zipcode level. Moreover, lower-level jurisdictions often override higher-level rates. Given that sales taxes are the largest source of state and local tax revenue, with a total revenue of \$500 billion in 2013, understanding how households respond to potentially non-salient changes in sales tax rates has important implications for both welfare and for local government revenue generation. Recent research on the effect of tax (non-)salience has therefore taken a particular interest in sales taxes.

We use the most detailed set of high-frequency state and local sales tax data available for the U.S. alongside household-level retail scanner data to test whether, and along which dimensions, consumers respond to these tax changes. We find that spending responds strongly to changes in sales tax rates but that spending on tax-exempt goods responds similarly to that of taxable goods. We show that this seemingly irrational behavior is fully consistent with rational forward-looking consumers in the presence of complementarities in ‘shopping costs’ across taxable and tax-exempt goods. Consumers can reduce the

¹Recent advances in modelling inattention and salience include ?, ?, ?, ?, ?, ?, and ?.

²A non-exhaustive list of empirical studies on tax salience, which we discuss in more detail below, includes ?, ?, ?, ?, ?, ?, and ?. Recent theoretical studies of the effect of tax non-salience on optimal taxation include ?, ?, and ?.

³Exceptions include excise taxes on goods like alcohol or gasoline, which are included in posted prices and generally are in addition to sales taxes, rather than replacing them.

need for future shopping trips in the short run by also stocking up on tax-exempt goods while shopping for taxable goods when tax rates are still low. We show that a standard life-cycle consumption model augmented with shopping complementarities across storable goods can quantitatively closely replicate the observed spending patterns of both taxable and tax-exempt goods.

We test this new mechanism along several dimensions, taking advantage of the detailed information on consumers and retailer locations. First, we directly test whether consumers cut down on shopping trips in the near future after stocking up on inventory in the month before a sales tax increase. We find that the number of store visits indeed decreases in the month after a sales tax increase, evidence of so-called multi-purpose shopping trips where consumers frequent multiple stores on the same trip thereby spreading the fixed cost of a trip over several store visits. Second, we exploit heterogeneity in “revealed costs” of shopping across consumers, reflected in the consumer’s average shopping frequency in the sample. Consumers that shop frequently in the absence of a tax change reveal that they face lower shopping costs than infrequent shoppers. Consistent with shopping cost complementarities, we find that infrequent shoppers indeed stock up more on tax-exempt goods than frequent shoppers who benefit less from such behavior. Third, we implement a form of a placebo test by looking at the relative response of taxable and tax-exempt online and catalog spending. Shopping complementarities should be absent or minimal when shopping online. We find that the responses of taxable and tax-exempt goods to a sales tax increase indeed differ substantially: online shopping for taxable goods significantly increases while online spending on tax-exempt good does not, consistent with the absence of shopping trip complementarities when shopping online.

Although some previous literature is supportive of the non-salience of sales taxes, we show that on average sales tax rate changes prompt large spending responses in the short-run. Using two strategies, we provide additional evidence of the salience of sales tax rate changes. First, we document several mechanisms that consumers exploit to avoid or mitigate the effect of the sales tax increases, all of which require taxes to be at least partially salient and for consumers to be forward-looking and optimizing. For instance, consumers bring spending forward to the months before a sales tax increase and more so for storable and durable goods. Consumers also engage in strategies that go beyond intertemporal substitution. Consumers that have the opportunity to shop across sales tax boundaries increasingly do so following a tax increase, and households also increase their online shopping. Second, we provide direct evidence of tax salience by estimating the dynamic response of Google searches containing the term “sales tax” by users in the affected tax jurisdictions around a sales tax change. We find a significant surge in Google

searches in the months prior to sales tax change that lasts a few months after the change, showing that consumers are actively acquiring relevant information before the change occurs. The response of Google searches also suggests that higher levels of attention are paid to sales taxes around the dates of the change.

While we find that salience of, and responses to, sales tax changes are high on average, we also find significant support for the importance of sales tax salience in governing the magnitude of household spending responses. We find that tax changes that receive relatively more newspaper coverage elicit larger responses (controlling for the size of the tax change). Similarly, spending responses are larger for changes in sales taxes that are driven by (highly-advertised) ballot propositions compared to legislated tax changes that are directly implemented by the state legislature.

This study therefore provides significant evidence that households generally respond rationally to changes in sales tax rates. Changes in sales tax rates seem to be highly salient, on average, and two realistic extensions of the standard life-cycle consumption model, storability and shopping trip complementarities, go a long way in explaining the observed household responses to sales tax changes. In particular, these extensions can fully explain the seemingly irrational response of tax-exempt spending.

This paper also highlights a methodological issue. Shopping complementarities affect the choice of empirical research design. In our case, using apparently unaffected tax-exempt goods as a control group in a difference-in-difference estimation would result in a substantial downward bias in our estimates that would completely mask the large responses to the sale tax changes. To deal with this issue, we instead take advantage of the high-frequency spending data and the predictive power of the model for short-run spending behavior around a sales tax change under the hypothesis of tax salience. Hence, our study highlights the fact that a difference-in-difference approach is not model-free, and its failure in this case is not due to general-equilibrium effects but holds even in partial equilibrium using within-household spending variation.

Related literature: This study touches on several strands of literature. It most directly relates to a recent surge in behavioral public finance research that studies the non-salience of tax rates and attributes, and studies the implications for optimal taxation. Empirical studies on this topic include the seminal contributions by ?, who find that households do not incorporate sales taxes into purchasing decisions and do not respond to changes in sales tax rates, and ?, who finds that how EZ tolls are paid, via transponder or via cash, affects driving behavior. ? is an early study of the effect of tax salience, demonstrating that consumers do not fully incorporate non-posted vehicle personal property taxes when

purchasing automobiles. ? and ? both find a differential household response to property taxes depending on whether the tax is paid separately or bundled with mortgage payments. ? report evidence for differences in household responses to equal-sized sales tax and income tax credits when purchasing hybrid automobiles. ? find evidence for low salience of sales taxes in an experimental setting with consumers making decisions about food consumption.

Recent theoretical work reviews the implications of non-salient tax attributes and rates for the theory of optimal taxation, extending the classic Ramsey, Pigou, and Mirrlees results—which are based fully optimizing agents—to settings that allow for “behavioral” agents who deviate from full optimization. See, for example, ?, ? and ?.

We also relate to the literature on the elasticities of intertemporal substitution of consumption and spending. However, our paper focuses heavily on the intertemporal substitution of spending rather than that of consumption. While some consumption changes surrounding these post-tax price changes are to be expected, the majority of the response that we observe is for durable and storable goods. This finding mirrors those of ??, and ?, who examine pre-announced VAT changes in New Zealand, Japan, and Germany.

Other work by ? and ? finds this effect on durables present in the United States when examining the impact of temporary monetary incentives to purchase new automobiles. Our work extends this research to a setting where the post-tax price is generally unobserved at the point of product selection and demonstrates a new extension, shopping complementarities, to household purchasing behavior that has strong implications for the estimation of elasticities of demand in the face of widespread price changes.

The rest of the paper is organized as follows. Section 1 describes the various data utilized in the paper. Section 2 estimates the response of taxable household spending to a pre-announced sales tax increase and provides a comprehensive study of the various mechanisms employed by consumers to avoid or mitigate those tax increases. We provide more direct evidence of tax salience captured by changes in Google searches for the term “sales tax.” This section also investigates the heterogeneity in spending responses as a function of tax salience, using newspaper coverage data and ballot initiatives. Section 3 estimates the response of tax-exempt spending to a sales tax increase and provides evidence of shopping complementarities. Section 4 provides extensions along several dimensions, including an analysis of the effect of tax salience on the sales tax elasticity of retail spending, and the identification of income and wealth effects from changes in news about a future sales tax rate changes. Section 5 extends the standard life-cycle consumption model to include shopping complementarity and storable-good inventories,

and calibrates it to match empirical moments derived in Sections 2 and 3. Section 6 concludes.

Data

We utilize a number of datasets to examine how sales taxes affect consumer spending.

Sales Tax Data

For data on local sales tax rates, we turn to Thomson Reuters OneSource sales tax service. This source allows us to construct a database of ZIP code level sales tax rates at a monthly frequency from 2008 to 2014 that covers the entirety of the United States. The data contain comprehensive information on all sales taxes imposed in a given ZIP code stemming from the state, county, city, and from special tax rate districts that the ZIP code is located in, such as school or water districts, police jurisdiction, etc. Moreover, there is information on the combined sales tax in a ZIP code, which may differ from the sum of all of the aforementioned sales tax rates due to statutory maximum sales taxes imposed at a state level (eg. state sales tax is 4% and the state imposes a maximum total local sales tax rate of 5%) or the fact that a lower-level tax jurisdiction such as a city overrides the sales tax rate of a higher-level jurisdiction, such as state sales tax rate. Our final sample includes over 41,000 ZIP codes from the 50 states and Washington DC.

Table 1 provides summary statistics regarding all state and local tax changes. Overall, sales tax changes in our sample period are highly asymmetrical. Over 84% of total observed changes in sales tax rates are positive, with average total sales taxes increasing from about 6.4% in 2008 to 6.9% in 2014. Restricting to changes in state sales taxes, we find that about 75% of changes are positive, with state sales taxes increasing from 5% to 5.5% on average over the decade to 2014.

State sales taxes generally make up the majority of total sales taxes in a given ZIP code. We therefore augment our sample with hand-collected state level changes in sales tax rates from 2004-2008 to match the sample period of the retail scanner data described below. State level changes also tend to be larger than local sales tax changes. For instance, the 25th and 75th percentiles of state level changes are 0.25% and 1%, while for all other local sales tax changes, they are 0.1% and 0.5%. Local changes are driven overwhelmingly by changes in city and county level taxes, while other sales taxes covering metro areas, water districts, school districts, or other geographic groupings play a much smaller role.

Retail Spending Data

Retail spending data is obtained from the Nielsen Consumer Panel Database at the Kilts Center (NCP, formerly the Homescan Consumer Panel). The NCP consists of a long-run panel of American households in 52 metropolitan areas from 2004 to 2014. The NCP aimed at measuring household demographic characteristics, household income, and spending on retail goods. Using bar-code scanners and diary entries, participants are asked to report all spending on household goods following each shopping trip. Monetary prizes and other drawings are utilized to incentivize higher levels of engagement.

The NCP is constructed to be a representative sample of the US population. Demographic survey information about participants is obtained when they join the panel as well as each year thereafter. Nielsen attempts to maintain a high quality of data with regular reminders to participant households that prompt them to report fully, and will remove non compliant households from their panel.⁴ ? provide a more detailed description of the NCP. ? perform a thorough analysis of the Homescan retail spending data, finding generally accurate coverage of household purchases though having some detectable errors in the imputed prices Nielsen uses for a subset of goods. Overall, they deem the Nielsen Homescan data to be of comparable quality to many other commonly-used self-reported consumer datasets.

Overall, there are more than 150,000 households in our sample. For the purposes of this paper, we choose to exclude households that change ZIP codes at any point in their time. This exclusion is done because we generally cannot tell the exact month of a move, so any change in sales taxes that accompany such a move may generate a spurious relationship with observed retail spending. Following these exclusions, over 135,000 households remain, yielding over 6 million household-month observations.

Given the nature of the data collection, Table 2 shows that the spending captured in the NCP consists primarily of trips to grocery stores, drugstores and pharmacies, as well as other mass-merchandise retailers. The types of goods purchased span groceries and drug products, small electronics and appliances, small home furnishings and garden equipment, kitchenware, and some soft goods.

In order to categorize individual products into taxable or tax-exempt goods, we first categorize products into one of several broad categories. We assign goods to one of the following categories: groceries, clothing, prepared food, medication, beer, liquor, wine, cigarettes, and non-exempt goods. We choose these categories to cover the range of categories that are treated differently on a state-by-state basis when it comes to determining whether a product is exempt from the sales tax. Groceries, for example, are almost always

⁴Approximately 80% of households are retained from year to year.

exempt from any state or local sales tax. However, the treatment of clothing or prepared food differs by state.

Our categorization is done at the Product Group level, as defined by the Nielsen Consumer Panel, so the types of goods which are ‘taxable’ or ‘tax-exempt’ will vary somewhat across states. There are 119 total Product Groups in the Nielsen data that we then assign to these 9 broader exemption categories. For instance, “Crackers”, “Dough Products”, “Fresh Meat”, and “Fresh Produce” would be Product Groups categorized as ‘grocery’ purchases. “Prepared Food Ready to Serve” is assigned to the ‘prepared food’ category, while “Soft Goods” are treated as ‘clothing’. A wide range of goods such as “Automotive” products, “Hardware and Tools”, and “Toys and Sporting Goods” are categorized as ‘taxable’ since goods of that type are not tax-exempt in any state in the United States.

Overall, the NCP tracks a sizable amount of a household’s spending on material goods. On average, we observe over \$350 of spending per month for each household. About half of this spending is on goods formally exempt from sales taxes while half is subject to sales taxes.

Tax Salience Measures

To obtain direct measures of tax salience we use newspaper article counts, Google searches and tax changes triggered by state-level ballot propositions.

Newspaper Article Counts: First, we employ data from the Access World News Newsbank database to measure news coverage of sales taxes at both a state and local level.⁵ We query a set of over 3,000 national, state, and local US newspapers at a monthly frequency from 2008 to 2016. Our query obtains the number of articles for each city-month or state-month that mention the term ‘sales tax’ or ‘sales taxes’. We exclude classified ads and restrict our search to newspapers rather than newswires or magazines. Raw counts of articles may give a misleading measure of news coverage of sales taxes given changes in the number and size of newspapers at any given time. To better gauge relative news coverage, we normalize each monthly value by the total number of newspaper articles written in that month and location.

We conduct searches at two levels of geographic aggregation. The first is at a state level (including Washington DC as its own state). The second is at a city level, where we attribute newspapers to cities based on Access World News’ categorization. Given that both our sales tax and retail spending data are at a ZIP code level, we match states and cities to ZIP codes using the city-state-ZIP matches in the Thomson Reuters sales tax

⁵ <http://www.newsbank.com/libraries/schools/solutions/us-international/access-world-news>.

data. This method yields a good match, with only 77 out of 1468 cities with newspapers being unable to be matched to ZIP codes in our sample.

Google Search: Second, we use Google search data obtained from Google Trends from 2004 to 2016 to study the search behavior of consumers around sales tax rate changes. Google Trends is a Google application that gives a time series of the relative amount of local search activity for specific search terms on Google.com.⁶ The values of Google Trends represent the number of searches on Google.com for the specified search term relative to the total number of searches on Google.com derived from a sample of all Google search data. Google Trends is normalized such that the highest value for the entire time period and term is set equal to 100. Its range of values is always between 0 and 100, where higher values correspond to higher ratios of total searches on Google.com for a given search term.

A potential concern, discussed in detail by ?, is that Google Trends imposes thresholds for reporting search data below which it imputes a zero value. For instance, too few searches were done for the search term ‘econometrics’ in July 2006 in Texas. Therefore, Google Trends displays a 0 rather than a low number, producing large swings in the time series data. For the term ‘sales tax’, there are a large number of zeroes between 2008 and 2010 in smaller states. We treat these values as missing data rather than true zeroes, due to the censoring that Google employs. In the years after 2010, there are only a few zeroes per year. Our results are robust to excluding all data from the years prior to 2011.

State Ballot Propositions: Third, using Ballotpedia.com we identify all state ballot propositions that involve changes in state sales taxes from 2004-2015. These data include propositions in Arizona, Arkansas, California, Colorado, Georgia, Maine, Massachusetts, Michigan, Minnesota, Missouri, South Dakota, and Washington, with some states having multiple ballots regarding sales taxes.

In total, we observe 20 propositions with potential effects ranging from a decline in sales taxes of 3.25% to an increase in sales taxes of 1%. 10 of the 20 propositions were successful, 9 were unsuccessful, and one was partially successful (took effect in a subset of state counties). 9 of the 20 propositions took place in November with the remaining propositions spread across February, May, June, and August.

Response of Taxable Spending to Sales Tax Changes

Our primary empirical methodology is to utilize a difference-in-differences specification. For most of our examination of the impact of changes in sales tax rates, we look at

⁶ <http://www.google.com/trends>.

monthly changes in spending at a household level. By construction, the control groups are those households who did not experience a change in the sales tax rate that they face in that month. For most regressions, both period and household-level fixed effects are incorporated, thus controlling for seasonal effects, macro effects, and allowing for household-level trends over time.

Taxable Retail Spending Response to a Sales Tax Increase

Table 3 shows how retail sales of goods subject to sales taxes, as measured by the Nielsen Consumer Panel, change following a change in the sales tax rate. We test the response of taxable goods to changes in total and state sales taxes, respectively, using a difference-in-differences approach,

$$\Delta \ln(x_{ht}^T) = \beta \Delta \tau_{ht} + \gamma_h + \eta_t + \lambda' z_{ht} + \epsilon_{ht}. \quad (1)$$

$\Delta \ln(x_{ht}^T)$ is the log-change in taxable expenditures by household h in month t , and $\Delta \tau_{ht}$ is the change in the sales tax rate in that month in the corresponding ZIP code, respectively in the corresponding state in Column 2. γ_h are household fixed effects and η_t are period fixed effects (year and month indicators). z_{ht} are additional time-varying co-variates at the level of the household or tax jurisdiction used in columns 3 to 8.

We restrict the main analysis to tax increases both because the vast majority of tax changes in our sample period are tax increases, and because the model below featuring storable and durable goods implies an asymmetric response to sales tax changes. Stocking up before a sales tax increase is easier and more likely synchronized across household in the month before the tax increase compared to letting inventories of storable and durable goods deplete in anticipation of a sales tax decrease. We analyze the robustness of our results to including all tax changes, including tax decreases, in Column 8 below.

Column 1 shows that following an increase in combined total sales tax rate of one percentage point (e.g., from 3% to 4%), taxable household retail spending decreases by 1.9%. This change in spending is measured in the month that the tax change occurs in relative to the month prior to the tax change.⁷ Column 2 restricts the analysis to state sales tax increases, which allows us to extend the analysis back to 2004, the start of the Nielsen Consumer Panel. The sales tax elasticity of taxable expenditures is almost identical to the one estimated using total sales tax changes in Column 1, although it is estimated with less precision due to the fewer tax changes (despite the longer sample period).

⁷The vast majority of sales tax changes go into effect on the first of the month, so the entire month is under the new sales tax rate. Our estimates are robust to excluding or weighting sales tax changes that occur on a different day of the month (the 15th is the second most common day).

Column 3 shows that controlling non-parametrically for time-varying household characteristics like income and family composition has little impact on the size of the sales tax elasticity of taxable spending. Column 4 adds in local and state unemployment rates to control for local business cycle conditions. We see little change in the coefficient of interest following the addition of these controls. Similarly, dropping the months from January 2008 to June 2009 that were part of the Great Recession according to the NBER recession dating committee also sees little change in the estimates, as seen in Column 5.

Column 6 employs an IV specification, instrumenting for changes in local sales taxes with changes in state sales taxes to better deal with measurement error. Again, we find similar relationships between changes in taxes and taxable retail spending. In Column 7 we utilize only within-state variation in sales taxes, including highly granular state-year and state-month fixed effects, and find that the magnitudes of our estimates remain virtually unchanged. Finally, Column 8 shows that the main results are robust to extending the analysis to include all tax changes, including tax decreases.

Overall, Table 3 shows that households strongly respond to changes in sales tax rates, which indicates that consumers are, on average, well aware of the changes.

Inspecting the Mechanism: Sales Tax Avoidance Strategies

In this section we inspect various mechanisms that might drive this large response. All of the mechanisms below require taxes to be at least partially salient and consumers to be forward-looking and optimizing. We then provide further direct evidence of sales tax salience using Google Searches and newspaper coverage of sales tax changes, and we study how tax changes that are presumably more salient to the consumer—either because they received more newspaper coverage or because they were initiated by a ballot proposition instead of being directly legislated—change the sensitivity of the spending response.

Jurisdictional Tax Arbitrage: ‘Cross-Border’ Shopping

One way to avoid paying more sales taxes is by engaging in cross-border shopping, taking advantage of lower rates in neighboring tax jurisdictions. To analyze this mechanism, we leverage one benefit of the Nielsen Consumer Panel, its ability to observe details of the shopping trips that households took including the type and location of a retailer. The NCP identifies stores by their three-digit ZIP code. In conjunction with the location of the household, this allows us to determine what fraction of household spending was conducted in an ‘alternate’ three-digit ZIP code (outside one’s ‘home’ three-digit ZIP code).

Column 1 of Table 4 tests whether this ratio responds to changes in local sales taxes,

finding no significant effect. However, it is generally difficult for most households to switch to shopping in a different three digit ZIP code given that the average three-digit ZIP code spans over 1,000 square miles. So, we might expect that households who are already able to conduct such shopping trips (e.g., those who might live or commute near a state or three-digit ZIP code boundary) might be more sensitive along this margin. In order to test this, we run the following specification:

$$\Delta \text{FracAltZIP}_{ht} = \beta \Delta \tau_{ht} + \delta \Delta \tau_{ht} \times \overline{\text{FracAltZIP}_h} + \gamma_h + \eta_t + \epsilon_{ht}. \quad (2)$$

FracAltZIP_{ht} is the fraction of household h 's total spending that is an alternative three-digit ZIP code outside its own home ZIP code (i.e., residential five-digit ZIP code), and $\overline{\text{FracAltZIP}_h}$ is the average over all household-months. Column 2 shows precisely this mechanism, a significant increase in alternate-ZIP spending for households who had already been conducting some of their shopping in alternative three-digit ZIP codes. This signals that, for households who could conceivably substitute spending into a different three-digit ZIP code, an increase in the sales tax made them shift additional spending to that ZIP code.

An important note about substitution across jurisdictions is that this pattern of behavior is evidence for strong impacts of sales tax changes on spending behavior, actual household consumption is affected to a much smaller degree. A recent study by ? also looks at the geographical substitution patterns surrounding sales taxes. They use credit card spending data to examine how ZIP code level spending is impacted by changes in sales taxes on both sides of the border of the tax jurisdiction, finding an elasticity of approximately 4.2 in ZIP codes that are located on state borders. Our results here align with their own. They also note persistent substitution to online retailers following sales tax increases.

Online Shopping

Another potential way for households to avoid increases in sales taxes is to shop online or via catalog and mail order (hereafter just referred to as ‘online spending’ or ‘online merchants’). Online merchants are generally not required to collect sales taxes for sales to purchasers if the merchant does not maintain a physical presence in the same state as the purchaser. During our sample period, a majority of online purchases were done without purchasers paying sales tax. Instead, households are officially required to pay a ‘use tax’ to their home state when completing their annual taxes. However, compliance with the use tax is extremely low. For instance, only 0.3% of California tax returns reported *any*

use tax related purchases in 2009. Because of this, households may shift purchases online where possible when local sales taxes increase.

Fortunately, the Nielsen Consumer Panel data categorizes purchases made from online merchants separately from brick-and-mortar retailers. In Column 3 of Table 4, we separately examine the impact of changes in sales taxes on these online purchases. We find that household shift spending to these online merchants in the month following a sales tax increase.⁸ These coefficients suggest that online spending in an affected household increases 1.6% following an increase in the sales tax rate of 1%.

Our estimates are consistent with recent estimates of the effect of taxation on online commerce. For instance, using state sales tax rate changes and purchase data from eBay, ? find an online-offline substitution elasticity of 1.8, which is in line with our estimate of 1.6. Similarly, ? estimate a tax elasticity of online purchases of -1.1 using Amazon’s staggered introduction of sales tax collection across different states in different months. This estimate is consistent with the ones reported above given that their experiment is a relative *increase* in the taxation of online purchases, while the previous experiments are relative *decreases* in the taxation of online purchases, i.e., an increase in the taxation of purchases from brick-and-mortar stores.

Intertemporal Substitution

Since sales taxes are almost always announced significantly in advance, we might expect that changes in household behavior precede the effective date of the sales tax change *if* households are forward-looking and sales tax rate changes are salient. To test this prediction, we estimate a dynamic version of equation (1) including leads and lags of the sales tax rate changes.

Figure 4(a) plots the patterns of spending in the three months prior to and the five months following any change in state level sales taxes and is scaled to reflect a 1% increase in sales taxes. We find elevated (depressed) levels of spending in the period preceding a sales tax increase (decrease) that quickly disappear once the change takes effect. For a sales tax increase of 1%, we see a dramatic fall in spending from the one month prior to the change to the month of the change, equivalent to a decline in spending of about 2.5%, similar to the more static responses reported in Table 3.

While this decline is steeper for taxable goods, exempt goods see a similar build-up prior to the tax increase and undergo a similar fall in the months afterwards. Importantly, we find that any significant deviation of levels of spending from their long-run averages

⁸We also find positive effects when looking at the fraction of spending done online rather than the change in dollar amounts.

tends to disappear after only a few months. That is, the short term response of household spending is significantly different than, and greater than, any long term response.

Storability and Intertemporal Substitution: We note that a large component of the change in retail sales in response to changes in sales taxes comes from intertemporal substitution within the household. When faced with an imminent sales tax increase, households bring forward purchases and spending to the lower-taxed period. Given this revealed desire to shift spending forwards in time, we would expect to see this substitution manifest itself to a larger degree for goods that are more durable and more storable. That is, it would not be feasible to purchase a several-month supply of milk (a tax-exempt good in most states) or carbonated beverages (taxable in most states) given that it would go bad before it could all be used.

To examine whether this pattern holds true empirically, we must first categorize all products in the Nielsen Consumer Panel data by their durability and storability. We do so in two ways.

First, we use average purchasing patterns in the data to inform us about the storability of each product. To do so, we categorize each product group with a continuous measure of how frequently products in a given group were purchased. For instance, milk, purchased weekly by an average household, would have a value of approximately $\frac{1}{4}$ (4 average purchases per month), while light bulbs may have a value greater than 2 (purchased less than once every two months). This ‘shopping cycle’, which is the purchase frequency of a typical product in product group g , is calculated first at a household level and then averaged across all households in the sample,

$$ShoppingCycle_g = \left(\frac{1}{N_H} \sum_{h=1}^H \frac{1}{T_h} \sum_{s=1}^{S_h} \mathbb{1}\{\text{Trip}_{sg}\} \right)^{-1}. \quad (3)$$

$\mathbb{1}\{\text{Trip}_{sg}\}$ denotes whether shopping trip $s \in S_h$ is one in which household h made a purchase from product group g . T_h measures the total number of months household h is in the sample and N_H represents the total number of households in the overall sample.

Second, we manually categorize the 119 product groups contained in the Nielsen data following previous literature on storable and durable goods; see e.g., ?. We construct three broad bins from this data, non-durable, semi-durable, and durable. These measures are meant to both capture how durable but also how storable a product is. For instance, a product like laundry detergent is not durable in the sense that it is consumed after use. However, it is highly storable and can be purchased well in advance of the use date.

Our rule of thumb is to categorize goods that need to be used within 6 months as non-

durable, those that can be kept up to 1 year as semi-durable, and those that can be kept or used for over 1 year as durable. We look at sources including Nielsen documentation, nutritional data on average shelf life, and data about average expiration dates for products to get a better idea of how durable and storable a given product group is. For instance, items like “Gum”, “Milk”, “Fresh Produce”, “Ice Cream”, and “Candy” are categorized as non-durable. “Beer”, “Canned Food” and “Coffee” are semi-durable, and “Soft Goods”, “Automotive”, “Laundry Supplies”, and “Toys” are labeled as durable.

The last column of Table 5 shows the product groups with the highest and lowest shopping cycles. While tax-exempt goods have on average a shorter shopping cycle as seen by the fact that of the first three product groups are tax-exempt in most states (“Bread and Baked Goods”, “Fresh Produce”, “Milk”) while “Carbonated Beverages” are typically taxable (and vice versa at the bottom, where flour is tax-exempt in most states and the other three products groups are typically taxable), there is a fair amount of overlap in storability across taxable and tax-exempt product groups.

The two measures of storability correspond fairly well with, with durable or storable goods, on average, having significantly lower purchases per month than non-durable or non-storable goods. In the durable category, the average purchases per month of a given product group is 0.38, for semi-durables, it is 0.8, and for non-durables, the average number of purchases per month is 1.5.

We conduct our analysis of heterogeneous behavior across categories of goods at a state-month level by estimating the following regression:

$$\Delta \ln(x_{gt}) = \beta \Delta \tau_{gt} + \delta \Delta \tau_{gt} \times \text{Storability}_g + \gamma_{state} + \gamma_g + \eta_t + \epsilon_{gt}, \quad (4)$$

Storability_g is the measure of storability of product group g , either the continuous average inverse purchase frequency or the discrete hand-classified storability indicator. Collapsing the data to the state-month level minimizes problems that arise from individual households having large numbers of zeroes for their monthly spending on particular fine categories, but also substantially reduces the number of observations and the power since we restrict the tax variation to state tax changes. Anticipating our findings in the next section, we do not distinguish between taxable and tax-exempt spending here in order to maintain sufficient statistical power. In the next section, we show that tax-exempt spending indeed responds similar to taxable spending, and we analyze the mechanisms behind this seemingly puzzling result.

Table 5 reports results from estimating Equation 4. In Column 1, we show the baseline spending response to a change in state sales taxes. On average, across categories, we see a decline in spending of approximately 2.3% in the month of a sales tax increase of 1%.

This state-level estimate corresponds well with our individual household-level responses of approximately 2% in Table 3. In Column 2 we interact the change in sales tax rates with our product group level data on inverted logged purchase frequency. We find that products purchased more frequently tend to be less affected by a change in sales taxes while infrequently purchased products see a larger than average response in the month of the tax change. To relate our new measure to the previous literature, Column 3 uses the simple indicator for durability as an interaction variable. Here we find a negative point estimate but it is not statistically different than zero.

Columns 4 to 6 include leads of changes in state level sales taxes to analyze the intertemporal substitution pattern documented in Figure 4(a) now by product storability and durability. That is, for the month of each state sales tax change, we observe the response of spending for the month prior to the change as well as the month after. In Columns 5 and 6, we find that not only do more durable and more storable products have larger declines in the month of a sales tax increase, but they see larger build-ups in the month prior to the increase. In fact, durable and storable products are the only product categories that see any increase in spending in the month before a sales tax increase.

These results are consistent with recent studies of intertemporal household spending behavior. ? for instance also finds that this pattern was seen around changes in the sales tax rate in New Zealand. Using data regarding three large changes in the national sales tax (Goods and Services Tax) rate, he finds strong evidence for intertemporal substitution among both durables and non-durables. Consistent with our estimates he finds that the magnitude of the substitution from the month of the change to the month prior to the change is 3 to 5 times larger for durable or storable goods than for non-durable and non-storable goods.

Direct Evidence of Tax Salience

The previous section shows a remarkable degree of tax foresight and forward-looking behavior on the parts of households, documenting how consumers evade sales taxes along several margins, all of which require sales tax changes to be salient to a large extent. In this section we provide more direct evidence of sales tax salience. First, using newspaper article ratios we document that the news media covers upcoming sales tax rate changes well in advance. Hence, households have easy access to relevant sale tax information well in advance of the tax changes. Second, and more importantly, using Google Search data we show that households *actively* acquire information about sales tax rates in advance of the tax rate changes. Both results are consistent with our findings that households adjust their spending patterns around sales tax changes.

Evidence from Newspaper Coverage

Figure 3 shows the two metrics of attention to and salience of sales taxes in the months surrounding state sales tax rate changes. The top panel displays the evolution of the ratio of news articles that mention sales taxes in the 10 months before and after a change in sales tax rates. It displays the regression coefficients of estimating a dynamic version of equation (1) with the log-level of the newspaper article ratio on the left-hand side and leads and lags of monthly state sales tax rates changes as the main dependent variable while controlling for state and time fixed effects. The dashed lines represent 95% confidence intervals using standard errors clustered at the state level. As noted in Section 1, the newspaper article ratio is derived from over 3,000 national, state, and local US newspapers aggregated at state levels. We find a gradual increase in articles, with the article ratio being significantly higher than the baseline level starting approximately 6 months prior to the change. In the month before the change occurs, the ratio peaks at a ratio about 75% higher than the baseline.

This figure is scaled by the size of the change so larger sales tax changes tend to get more news coverage relative to smaller changes. Following the change, news about sales taxes quickly recedes to the baseline level, with the ratio being statistically indistinguishable from zero just one month following the change.

Evidence from Google Searches

While the top panel shows that in households have in principle readily access to the latest information about upcoming sales tax changes, it remains to show that households actively acquire this information. The bottom panel of Figure 3 shows that households indeed increase the search for information about sales taxes in anticipation of a sales tax change. The figure plots the coefficients of the same specification as before, replacing the newspaper article ratio with Google Search index. As with the newspaper-based measure, search peaks in the month before a change takes place, rising to over 130% of the baseline level of search. Google searches about sales taxes do not respond as far in advance of the change occurring, but have significantly elevated levels for a longer period of time than does the news-based measure. This may reflect a subset of households only realizing sales taxes may have changed over a longer period of time.

Response of Tax-Exempt Spending

In theory, we might expect that the effect of a sales tax increase on tax-exempt spending would be zero, or would capture any income or wealth effects of the tax reform. Hence,

it seems that tax-exempt spending could be used as an ideal control ‘group’ to estimate the effect of a tax change using a within-household difference-in-differences specification. However, there are a few reasons why we might still see an effect even for goods that are not directly affected by sales tax rate changes. Households may be unaware of the fact that some goods are exempt from sales taxes or may simply mis-attribute an exempt product to a non-exempt category. After all, the laws defining which goods are exempt and non-exempt are quite detailed and technical, so it would not be surprising to have a significantly level of this sort of error on the part of households.

In addition, sales tax changes apply to many goods at the same time and hence are different from idiosyncratic (pre-tax) price changes typically used to estimate demand. In addition to income effects mentioned above, sales tax changes can thus also affect store traffic in the same way deals or large store sales are used by retailers for the purpose for cross-selling other goods with higher margins that are not on sale. In this section we therefore analyze the response of tax-exempt spending in order to assess the total effect of a sales tax rate increase on retail spending.

Reduced-Form Evidence

Table 6 reports the response of tax-exempt spending in the month of a sales tax rate increase relative to the previous month, mirroring the exact specifications seen in Table 3. We find a highly significant response of exempt spending to changes in both state and local sales taxes. This effect remains when instrumenting for local sales taxes with state sales taxes and when controlling for local business cycle conditions. Our estimates for tax-exempt goods are lower than for taxable goods, but both are significantly different than zero and are statistically indistinguishable from one another. Figure 4(b) shows that tax-exempt spending mimics the behavior of taxable spending not only in the month of the sales tax increase but also in the months surrounding that tax change.

While the impacts from changes in sales tax changes are similar to those found for taxable goods, one deviation from the taxable spending responses is the positive sign on the state level unemployment rates. Unlike spending on taxable goods, we find that spending on exempt goods tends to increase when the unemployment rate increases. This is likely because tax-exempt goods are made up primarily of grocery items, which tend to receive additional spending at the expense of restaurant spending during downturns and falls in household income (see, for example, ?). This effect may also be what drives the main coefficient further negative when we exclude NBER recession quarters in Column 5.

Inspecting the Mechanism: Shopping Complementarities

The reduced-form evidence shown in Table 6 and Figure 4(b) is consistent either with non-salience of tax-exemption status (i.e., consumer confusion) or with complementarities between taxable and tax-exempt spending arising from the short-run shopping cost savings achieved by also stockpiling tax-exempt goods while shopping for taxable goods before the tax increase. For instance, on a typical trip to a grocery store, a household may purchase both exempt and non-exempt goods (fresh produce, cookware, and a deli sandwich, for example). If households adjust purchasing responses to sales tax changes at a trip level, then we may expect that behavior of exempt and non-exempt goods would be correlated. Moreover, households with heterogeneous shopping costs would be predicted to have different spending responses for exempt goods.

In this section we provide new evidence for the role of shopping trip complementarities in explaining the observed patterns. We document this novel mechanism along three dimensions, taking advantage of the detailed information on consumers and retailer locations.

First, we directly test whether consumers save on shopping trips in the short run. Second, we exploit heterogeneity in “revealed costs” of shopping across consumers, reflected in the consumer’s average shopping frequency in the sample. Consumers that shop infrequently in the absence of a tax change reveal that they face higher shopping costs than frequent shoppers and hence should increase tax-exempt spending relatively more. Finally, we implement a form of a placebo test by looking at the relative response of taxable and tax-exempt online spending, since shopping complementarities should be absent or minimal when shopping online.

Table 7 presents the results of this analysis. In Columns 1 and 2, we directly test whether shopping trips themselves respond to changes in sales taxes or households simply purchase less per trip. We find strong evidence that the number of distinct trips (which may include visits to multiple stores) and the number of distinct store visits respond negatively to changes in sales tax rates. In particular, the number of trips falls by approximately the same magnitude as overall spending, suggesting that the trips adjustment margin is a dominant one.

Columns 3 to 6 investigate whether households with different ‘revealed shopping costs’ behave differently following a change in sales taxes. We first calculate the average number of shopping trips they make in a month for each household. We then assign the top 25% of households (with more than 19 trips per month) as ‘low-shopping-cost’ households and the bottom 25% of households (with fewer than 9 trips per month) as the ‘high-shopping-cost’ households. We propose that the average number of shopping trips a household takes

per month correlates negatively with the total costs of the trip, including transportation costs, inventory costs, and time costs, in line with the model in Section 5.

Columns 3 and 4 estimate the spending response of exempt and taxable goods for households with low shopping costs following a sales tax increase. We find no impact on exempt spending, while taxable spending declines approximately 2%. In contrast, for households we deem to be high-cost shoppers, both exempt and taxable spending fall nearly identically, suggesting that these households bundle their purchases to minimize the number of shopping trips that they must undertake.

Finally, Columns 7 and 8 test for the asymmetric response of online spending to sales tax changes as predicted by models with shopping trip costs. When shopping online, there are fewer gains to bundling multiple purchases at once, since no transportation costs need be incurred across different websites and online purchases are often made of single goods rather than a cart full of goods. Just as with the low-cost shoppers, we find that, following an increase in sales taxes, spending on exempt goods from online merchants is largely unaffected, while spending on taxable goods increases significantly. As with the results reported in Table 4, this response suggests that consumers evade sales taxes by substituting to online platforms and failing to pay use taxes on those purchases.

Overall, these results suggest that shopping complementarities play an important role in affecting the purchasing decisions of households. It also demonstrates the caution one must take when estimating price elasticities in a difference-in-differences framework, even in the absence of general equilibrium effects. Despite the fact that some goods' prices are unaffected, demand for them may shift due to changes in shopping behavior. This is true in our setting with tax changes, but also may be true when stores put portions of their goods on sale, or an appreciable number of items at a store undergo a price change at the same time.

Extensions and Robustness

In this section, we extend our analysis and test the robustness of the finding that changes in sales tax rates have large effects on household spending in the short run.

The Role of Tax Salience: Evidence from Ballot Initiative

A natural question that arises given the results displayed in Figure 3 is whether tax salience plays any role in the response to tax rate changes. While the results in Sections 2 and 3 document a significant degree of tax salience and foresight on average, it seems reasonable that some households are not fully aware of the tax rate changes, or some aspects of the tax code such as the exemption status of certain goods (e.g., cookies vs. candies)

is not fully salient. In this section we therefore test whether more salient tax changes elicit larger spending responses. This analysis is motivated by several highly influential previous studies that document a large degree of non-salience of sales tax rates among consumers (see the literature mentioned in the introduction).

In Table 8, we include two measures of salience and examine their impact on changes in household spending. The first is the aforementioned index of sales tax news coverage in the month prior to the change. Given that the size of the sales tax change strongly impacts the level of coverage, we first obtain the residuals from a regression of the amount of sales tax news coverage on the size of the change, the squared size of the change, and time fixed effects.

With this approach, we interpret the resulting residuals as a measure of news coverage of the impending sales tax change that is unrelated to the size of the change (ideally driven by the amount of other important news in that period, editorial decisions, etc.). Here, the assumption is that the more that sales taxes are written about in local newspapers, the more likely it is that a given household will be aware of the upcoming change in sales taxes and that they will be in position to react to the change.

Columns 1 to 3 of Table 8 interact this news-based measure with changes in state sales taxes. We again find that, in general, sales tax changes have a negative relationship with spending. Moreover, while not always statistically significant, changes that had more news coverage (conditional on the size of the change) also had larger declines, and this effect of news coverage strengthened as the sales tax change got larger. The coefficient on the interaction term of Column 1 shows that an increase in news coverage of 1 percentage point (e.g., 1.5% of news articles mention sales taxes rather than 0.5%) would increase the response to a 1% sales tax by about 50% (from -1.6% to -2.4%).

Columns 4 to 6 take a different approach to testing heterogeneity in household responses across sales tax changes. Here we utilize data on state-level ballot measures that changed state sales tax levels. Our prior is that sales taxes enacted through state-wide ballots would garner more media attention than those enacted through a vote solely by their state representatives and also would force all voters to be at least somewhat aware of the initiative that they are voting on.

We find that changes in sales tax rates that were authorized by a state-wide ballot measure tended to produce much larger responses among households. In fact, for the modal state level change (a 1% increase in sales taxes), we find that having the change come via the ballot box rather than the legislature increases the response of total household spending by approximately 150% (i.e., from -1.5% to -4%).

‘News’ Response: Income, Wealth and Substitution Effects

Columns 7 and 8 of Table 8 demonstrate some evidence for an announcement effect of sales taxes. For most of the changes in our sample, we are unable to determine when exactly the sales tax change was finalized (often 3 months to 12 months prior to the change taking place). For ballot provisions, however, we can precisely measure this date, allowing us to look for changes in household spending behavior prior to the change actually taking place.⁹

In a model with fully informed and rational consumers, households would perceive this future tax increase as an (permanent) increase in future prices. As documented by the model in Section 5, this leads to a spending response that is the combination of a negative income effect (the same pre-tax consumption plan is more expensive) and a positive intertemporal substitution effect (spending is temporarily cheaper in the period before the sales tax increase). In addition, there could be wealth effects that depend on the consumer’s perception and valuation of what the government plans to do with the additional revenue.

Column 7 demonstrates that this effect seems to be true, on average, across all ballots (whether they passed or failed), with the act of voting on the ballot being associated with a 0.6% decline in household retail spending. However, we can also separate these ballots into those that failed and those that passed. Here we find opposite signed coefficients that are statistically different from each other. Judging the point estimates, we find a near zero effect on spending following a failed tax increase initiative, while we see a much larger decrease in spending following a successful tax increase vote. This again speaks to the rational response of households with a high degree of fiscal foresight.

Robustness: Price and Quantity Response

While the majority of the paper discusses household responses in terms of changes in dollars of retail spending, it is possible that this may portray an incomplete view of household behavior. If retailers adjust prices or households shift spending to different types of goods, we may have a different interpretation of how the pre-tax spending response relates to actual consumption.

Table 9 examines two other important margins. Columns 1 to 4 mirror the analysis done in Table 3 and 6 but using log-changes of quantities (items) purchased as the dependent variable rather than log-changes in spending. We find qualitatively similar effects,

⁹Ideally we would weight the responses by how close the outcome of the ballot proposition was in order to interpret the spending response as a rational response to a change in expected tax rates. Unfortunately, we do not yet have this data.

with declines in quantities mirroring the declines in spending following an increase in sales tax rates. This indicates that households are likely not simply substituting lower quality and lower priced goods to reduce pre-tax spending.

Columns 5 to 8 test another potential confounding margin of adjustment. If retailers fully offset sales taxes, we might observe a decline in spending with no actual decline in consumption. Using data on both retail and wholesale prices we find that there is limited amounts of offsetting behavior on the part of firms. Retail prices decline by 0.15%-0.20% in the month following a 1% increase in sales tax rates, while wholesale prices remain unaffected.¹⁰

The non-response of the wholesale prices may be driven by the fact that wholesalers are less geographically concentrated and so do not price to local conditions to the extent that retailers do. Another reason may be that wholesale prices tend to be more stable and feature fewer short-term sales than do retailers, leading to somewhat higher menu costs and a reduction in desire to change prices at high frequency. We leave a more complete analysis of the firm response to sales tax changes, taking into account the response of demand documented in this paper, to future research.

Modelling the Spending Response to Sales Taxes

To guide our empirical analysis and to interpret our results we use a linear-homogeneous demand system for different types of potentially storable (or durable) goods as our benchmark model. The consumer's problem is to

$$\max_{a_t, c_{it}, s_{it}} \sum_{t=0}^{\infty} e^{-\rho t} u(C_t) \quad \text{with } u(C_t) = \frac{C_t^{1-1/\sigma}}{1-1/\sigma} \quad (5)$$

$$\text{s.t. } s_{it} = \delta_i s_{i,t-1} - c_{it} + x_{it} \quad (6)$$

$$a_t = e^r a_{t-1} + y - \sum_i p_{it} x_{it} - \theta(s_t) \quad (7)$$

$$(x_{it}, c_{it}, s_{it}) \geq 0 \quad (8)$$

with CES aggregator $C_t = \left(\sum_i b_i^{\frac{1}{\epsilon}} c_{it}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$ and ideal utility-based cost-of-living index $P_t = \left(\sum_i b_i p_{it}^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}}$, which is the minimum expenditure necessary to obtain one unit of the composite good C_t . Therefore, total expenditures on goods consumed (but not necessarily on all goods purchased) in period t is $P_t C_t$.

We distinguish between consumption c_i and expenditures x_i . Consumption of storable goods can deviate from expenditures on storables due to accumulation or decumulation of

¹⁰The pricing data are described in Online Appendix A. to save space.

inventory s_i . δ_i measures storability as the fraction of the inventory of good i that did not perish from the previous period and could be consumed in this period.¹¹ p_{it} is the price of storable good i in period t , r is the continuous-compounding net return on the financial asset between period $t - 1$ and t and y is non-financial income. $\theta(s_t)$ denotes inventory costs and benefits in addition to the opportunity costs of capital implicit in $p_{it}x_{it}$, such as storage costs and reductions in the number of shopping trips. θ typically depends on the inventory of all goods, $s_t = (s_{1t}, s_{2t}, \dots)$. For now we assume the function to be twice differentiable and to be convex in each single argument, i.e., $\theta''_{it} \equiv \frac{\partial^2 \theta(s)}{\partial s_{it}^2} > 0$ for $s_{it} \geq 0$. This condition guarantees that the decision between investing in financial asset a or in inventories s has a solution even if the marginal return on holding inventory temporarily exceeds the return on the financial asset. Hence, this assumption is necessary because a and s_i are both riskless assets.

An important assumption that makes this problem tractable is that the inventory cost function is independent of current consumption, $\frac{\partial \theta}{\partial c_{it}} = 0$. Hence, the consumption optimization problem and the problem of optimal inventory management are separable.¹² We also define the cross-term $\theta''_{ijt} = \frac{\partial^2 \theta(s)}{\partial s_{it} \partial s_{jt}}$, which captures whether inventories are complements or substitutes in terms of investments, but not necessarily for consumption (which is governed by ϵ). Finally, equation (8) assumes no reselling of storable goods. Hence, these are irreversible investments. For the goods in our scanner data set, which exclude large durables such as furniture or cars this is a reasonable assumption.

Hicksian compensated demand for good i depends on the good's real price p_{it}/P_t , holding utility C_t constant,

$$c_{it} = b_i \left(\frac{p_{it}}{P_t} \right)^{-\epsilon} C_t. \quad (9)$$

The consumption budget share of good i is $B_{it} \equiv \frac{p_{it}c_{it}}{P_t C_t} = \frac{\partial \ln P_t}{\partial \ln p_{it}} = b_i \left(\frac{p_{it}}{P_t} \right)^{1-\epsilon}$. Rewriting the intertemporal budget constraint (7) as $a_t = e^r a_{t-1} + y - \sum_i p_{it}(s_{it} - \delta s_{i,t-1}) - \theta(s_t) - P_t C_t$, the first-order conditions are

$$C_t = \left(\frac{P_t}{P_{t-1}} \right)^{-\sigma} e^{\sigma(r-\rho)} C_{t-1} \quad (10)$$

$$\theta'_{it}(s_t) = e^{-r} \delta_i p_{i,t+1} - p_{it} \equiv PV_{it}. \quad (11)$$

¹¹ We provide more details about the model in the appendix, where we also show that durables can be mapped into this storable goods framework.

¹² See ? and ? for a discussion of some of the issues that arise in the general case. Below we will assume that inventory costs depend on the (new) steady state level of consumption, \tilde{c}_i , instead of the current level of consumption, c_{it} . Similarly, we will model potential inventory cost complementarity as a result of (additional) synchronization of temporary out-of-steady-state shopping trips, which are the result of an anticipated sales tax increase.

Equation (11) holds only for inventory goods, i.e., if the optimal target level of inventory is positive, $s_{it} > 0$. In this case (11) equates the marginal benefits of holding inventory with the present value of the cost of the investment. Since θ is convex in s_i , optimal inventories are increasing in next period's price $p_{i,t+1}$ and the good's storability δ_i , and decreasing in the current price p_{it} and the opportunity cost of time, r .

Steady state: To simplify notation we assume $r = \rho$, which implies that the consumption index is constant, $C_t = C$. Using Euler equation (10) and imposing the no-Ponzi condition $\lim_{t \rightarrow \infty} e^{-rt} a_t = 0$ we can solve the budget constraint forward,

$$C \times P_0^\sigma \sum_{t=0}^{\infty} e^{-rt} P_t^{1-\sigma} \equiv C \times D = W, \quad (12)$$

where $W = e^r a_{-1} + \sum_{t=0}^{\infty} e^{-rt} \hat{y}_t$ is total wealth, with income net of total inventory costs $\hat{y}_t \equiv y - [\sum_i p_{it}(s_{it} - \delta_i s_{i,t-1}) + \theta(s_t)]$. In the steady state, taxes and all prices are constant, such that $p_{it} = p_i$ and $P_t = P$. Denoting steady-state values with a tilde, we obtain $\tilde{C} = (1 - e^{-r}) W/P$, which is the perpetuity value of real wealth, and $\tilde{c}_i = b_i \left(\frac{p_i}{P}\right)^{-\epsilon} \tilde{C}$.

Comparative statics

To analyze the consumption response, we focus on two separate composite goods in addition to the main composite good C_t , one that consists of goods that are subject to sales taxes, $c_{\tau,t}$ with price $p_{\tau,t} = (1 + \tau)p_{\tau,t}^{pre}$, and the other consists of tax-exempt goods $c_{e,t}$ with price $p_{e,t} = p_{e,t}^{pre}$. p_{it}^{pre} are posted pre-tax prices. The consumption index is now $C_t = \left[b_e c_{e,t}^{\frac{\epsilon-1}{\epsilon}} + b_\tau c_{\tau,t}^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}}$.

Suppose a consumer learns today ($t=0$) about a permanent sales tax increase in period $T \geq 0$ such that

$$d \ln(1 + \tau) \equiv \begin{cases} d \ln(1 + \tau_t) & \text{if } t \geq T, \\ 0 & \text{if } t < T. \end{cases} \quad (13)$$

In general $\frac{\partial \ln p_{is}}{\partial \ln(1 + \tau_t)} = \mathbb{1}_{(i,s)=(\tau,t)} + \frac{\partial \ln p_{is}^{pre}}{\partial \ln(1 + \tau_t)}$, where $\mathbb{1}_{(i,s)=(\tau,t)}$ is one if $i=\tau$ and $s=t$, and 0 otherwise. For now we assume $\frac{\partial \ln p_{is}^{pre}}{\partial \ln(1 + \tau_t)} = 0$. That is, we assume that pre-tax prices do not change in anticipation of the sales tax increase, and that there is full and immediate pass through of sales taxes on tax-inclusive prices, but no effect on prices of tax-exempt goods. In the empirical analysis we estimate the dynamic response of pre-tax retail prices to assess these assumptions.

Consumption response to a sales tax change: Taking logs of the steady-state consumption index level (12) and differentiating with respect to (future) sales taxes yields

$$\frac{d \ln C_0}{d \ln(1 + \tau)} = e^{-rT} \left(\underbrace{\sigma B_\tau \mathbb{1}_{T>0}}_{\text{intertemporal substitution}} - \underbrace{(B_\tau - \varepsilon_{\hat{y},\tau})}_{\text{income net of wealth effect}} \right). \quad (14)$$

The negative income effect is given by $-B_\tau$ while $\varepsilon_{\hat{y},\tau} = \frac{\partial \ln \hat{y}_t}{\partial \ln(1+\tau_t)}$ captures wealth effects of a sales tax change on net-of-inventory income \hat{y} . These wealth effects also include any off-setting taxes or government transfers. Below we estimate the correlation of sales tax changes with other local taxes to calibrate $\varepsilon_{\hat{y},\tau}$. The full size of the wealth effect depends on what the consumer expects the government to do with the additional sales tax revenue, which is outside of our model.¹³ The indicator function $\mathbb{1}_{T>0}$ accounts for the fact that there is no substitution effect σB_τ if a tax increase is unanticipated. In this case the consumption index falls by the combined income and wealth effect. On the other hand, if sales taxes are expected to change only in the distant future ($T \rightarrow \infty$), then there is no effect on current consumption.

For an anticipated tax increase ($T > 0$), we use Euler equation (10) to compute the (compensated) decrease in the consumption index when taxes increase at time t relative to the consumption level right after the announcement, C_0 ,

$$\frac{d \ln \frac{C_t}{C_0}}{d \ln(1 + \tau)} = -\sigma B_\tau \mathbb{1}_{t \geq T > 0}. \quad (15)$$

At the time taxes actually change (either anticipated or unanticipated), consumers also substitute away from taxable to tax-exempt goods according to Hicksian demand (9). Hence, the comparative statics of the effect of a permanent sales tax increase in period T on consumption in period $t \geq 0$ relative to the initial steady state $c_{i,-1} = \tilde{c}_i$ can be summarized as

$$\frac{d \ln c_{it}}{d \ln(1 + \tau)} = \begin{cases} -\left[(B_\tau - \varepsilon_{\hat{y},\tau}) - \sigma B_\tau\right] e^{-rT} & \text{if } t < T > 0, \\ -\left[(B_\tau - \varepsilon_{\hat{y},\tau}) - \sigma B_\tau\right] e^{-rT} - \sigma B_\tau - \epsilon(\mathbb{1}_{i=\tau} - B_\tau) & \text{if } t \geq T > 0, \\ -(B_\tau - \varepsilon_{\hat{y},\tau}) - \epsilon(\mathbb{1}_{i=\tau} - B_\tau) & \text{if } t = T = 0, \end{cases} \quad (16)$$

¹³ $\varepsilon_{\hat{y},\tau}$ also includes any mechanical effects of changes in inventory costs if the particular inventory cost function θ depends on steady-state consumption, as it does in our specification. To obtain closed-form solutions, we maintain the separability of the consumption and the inventory choice problem by assuming that the consumer ignores the effect of the consumption choice on inventory costs; see the appendix for more details.

with $i \in \{e, \tau\}$ and $\frac{d \ln c_{i,t}}{d \ln(1+\tau)} = 0$ for $t < 0$. The second line—the effect of an anticipated sales tax change on consumption starting in the period of the tax change—summarizes the different effects of the policy. The first two terms in square brackets are the negative income effect ($-e^{-rT} B_\tau$), which depends on the consumer's budget share of taxable goods, and the wealth effect ($e^{-rT} \varepsilon_{\hat{y}, \tau}$). The third term ($e^{-rT} \sigma B_\tau$) is the positive *intertemporal* substitution effect of consumption before period T . The consumer's ability to move consumption forward to periods with lower tax rates mitigates the negative net income effect. Since these effects all take place starting in period T , they need to be discounted back. The fourth term ($-\sigma B_\tau$) is the reverse of the the third effect, reflecting the negative intertemporal substitution effect on overall consumption after taxes are permanently increased in period T . Finally, the fifth and the sixth term are the *intratemporal* substitution effect of the change in relative prices starting in period T , which is positive for exempt goods, ϵB_τ , and negative for taxable goods, $-\epsilon(1 - B_\tau)$.

Inventory response to a sales tax change: The effect of a sales tax increase on the present value of a dollar invested in the taxable inventory good before total inventory costs, the right-hand side of (11), is

$$\frac{dPV_{\tau,t}}{d \ln(1+\tau)} = \begin{cases} -p_\tau(1 - \delta_\tau e^{-r}) < 0 & \text{if } t \geq T \\ p_\tau \delta_\tau e^{-r} \geq 0 & \text{if } t = T - 1 \\ 0 & \text{if } t < T - 1 \end{cases} \quad (17)$$

for an anticipated sales tax increase ($T > 0$), and $\frac{dPV_{\tau,t}}{d \ln(1+\tau)} = p_\tau(e^{-r} \delta_\tau - 1) < 0$ for an unanticipated tax change ($T = 0$), while the present value of the exempt inventory is not affected,

$$\frac{dPV_{e,t}}{d \ln(1+\tau)} = 0 \quad \forall t, T. \quad (18)$$

While there are no direct effects of sales tax changes on target inventories of tax-exempt goods as seen in (18), there can be indirect effects through the complementarity in total inventory costs, which are captured by the cross-derivative $\theta''_{\tau,e}$. For instance, consumers can save on future shopping trips by simultaneously also stocking up on tax-exempt goods so that $\theta''_{\tau,e} < 0$.

We use an augmented Allais-Baumol-Tobin model to derive a functional form of inventory cost function θ , which captures two sources of costs in addition to the holding costs implicit in $p_{it}x_{it}$.¹⁴ First, economies of scale from a reduction of shopping costs θ_s by

¹⁴ An alternative approach would be to assume transaction fixed costs, such as search and shopping trip

stocking up on goods that have similar steady-state inventory cycles due to similar storability or consumption demand for example. The costs are linear in the fraction of trips per period, which are approximately equal to $\frac{\tilde{c}_i}{s_{it}}$. Second, we augment the model to capture the additional cost complementarity θ_c due to economies of scope from synchronizing additional out-of-steady-state shopping trips,¹⁵

$$\theta(s_t) = \theta_s \sum_{i \in \{\tau, e\}} \frac{\tilde{c}_i}{s_{it}} - \theta_c \prod_{i \in \{\tau, e\}} \frac{\tilde{c}_i}{\tilde{s}_i} \left(\frac{s_{it} - \tilde{s}_i}{\tilde{s}_i} \right), \quad i \in \{e, \tau\}. \quad (19)$$

Steady state inventory $\tilde{s}_i = \sqrt{\frac{\theta_s \tilde{c}_i}{p_i(1 - \delta_i e^{-r})}}$ has the standard Baumol-Tobin square-root form.

Given this functional form, inventories are positive in steady state. Therefore, we can totally differentiate (11) around the steady state with respect to a sales tax change to get the inventory responses in period t of taxable and tax-exempt spending to an anticipated sales tax increase at date T ,

$$\frac{d \ln s_{\tau, t}}{d \ln(1 + \tau)} = \begin{cases} \psi (1 - \delta_\tau) < 0 & \text{if } t \geq T \text{ or } T = 0, \\ \psi \delta_\tau e^{-r} \geq 0 & \text{if } t = T - 1 \geq 0, \\ 0 & \text{otherwise,} \end{cases} \quad (20)$$

$$\frac{d \ln s_{e, t}}{d \ln s_{\tau, t}} = \phi \geq 0, \quad (21)$$

with inventory multipliers $\phi = \frac{\theta_c}{2\theta_s} \frac{\tilde{c}_\tau}{\tilde{s}_\tau}$ and $\psi = \left[1 - \left(\frac{\theta_c}{2\theta_s} \right)^2 \frac{\tilde{c}_e}{\tilde{s}_e} \frac{\tilde{c}_\tau}{\tilde{s}_\tau} \right]^{-1} \frac{1}{2} \frac{1}{1 - \delta_\tau e^{-r}} > 0$.

The dynamic response of observable expenditures

While neither consumption nor inventories are observable, we do observe expenditures.¹⁶ Hence, in order to map the model to the data we need to (i) relate the comparative statics (e.g., $d \ln s_{it}$) to observed dynamics (e.g., $\Delta \ln s_{it}$) and (ii) relate changes in consumption and inventories to changes in observed expenditures. We achieve the latter

costs. Unfortunately, optimal investment with fixed plus linear transaction costs and multiple inventory goods becomes a hard combinatorial problem; see ? for example.

¹⁵ The complementarity component uses a first-order approximation of the shopping-trip frequency around the steady state in a standard Baumol-Tobin inventory model. While the results are not affected by taking this approximation, it proves convenient for deriving a closed-form solution of the inventory responses.

¹⁶ Information on inventory holdings is rare even for firm-level data and much more so at the household level. Similarly, household consumption is typically not observable. For instance, both types of consumer surveys conducted by the Bureau of Labor Statistics (BLS), the Interview and the Diary portion of the Consumer Expenditure Survey, ask about expenditures, not about consumption.

by log-linearizing the law of motion of inventories (6) around the steady state,

$$\Delta \ln x_{it} \approx \frac{\tilde{s}_i}{\tilde{x}_i} \left[\Delta \ln s_{it} - \delta_i \Delta \ln s_{i,t-1} + \frac{\tilde{c}_i}{\tilde{s}_i} \Delta \ln c_{it} \right], \quad (22)$$

where Δ denotes temporal changes in contrast to comparative static changes d , i.e., $\Delta x_{it} \equiv x_{it} - x_{i,t-1}$, and $\frac{\tilde{s}_i}{\tilde{x}_i} = (1 - \delta_i + \frac{\tilde{c}_i}{\tilde{s}_i})^{-1}$.

The comparative statics of consumption in (16), (20) and (21) are all comparisons relative to the initial steady state, $c_{i,-1}$ and $s_{i,-1}$ respectively. Hence, temporal changes are differences in these comparative statics, e.g. $\Delta \ln s_{it} = d \ln s_{it} - d \ln s_{i,t-1}$. Therefore,

$$\frac{\Delta \ln c_{it}}{\Delta \tau} = \begin{cases} -[(1 - \sigma)B_\tau - \varepsilon_{\hat{y},\tau}]e^{-rT} & \text{if } t = 0 \text{ and } T > 0, \\ -\epsilon \mathbb{1}_{i=\tau} - (\sigma - \epsilon)B_\tau & \text{if } t = T \text{ and } T > 0, \\ -[(1 - \epsilon)B_\tau - \varepsilon_{\hat{y},\tau}] - \epsilon \mathbb{1}_{i=\tau} & \text{if } t = T = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (23)$$

for $i \in \{\tau, e\}$ and

$$\frac{\Delta \ln s_{\tau,t}}{\Delta \tau} = \begin{cases} \psi e^{-r} \delta_\tau \geq 0 & \text{if } t = T - 1 \text{ and } T > 0, \\ -\psi < 0 & \text{if } t = T \text{ and } T > 0, \\ -\psi (1 - e^{-r} \delta_\tau) < 0 & \text{if } t = T = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (24)$$

$$\Delta \ln s_{e,t} = \phi \Delta \ln s_{\tau,t}. \quad (25)$$

Simulating the model

To assess the quantitative predictions of the model, we calibrate it using empirical moments from the data as well as consensus estimates from the literature whenever possible. We use value-weighted average monthly purchase frequencies of taxable and tax-exempt expenditures to calibrate the ratios of consumption to inventories, $\frac{\tilde{c}_\tau}{\tilde{s}_\tau} = 1.1$ and $\frac{\tilde{c}_e}{\tilde{s}_e} = 1.6$, consistent with our micro-foundation of the inventory cost function. In the baseline specification we assume that the ratio of “own” to cross shopping trip costs is one, $\frac{\theta_c}{\theta_s} = 1$, since both parameters capture the same opportunity cost of time. We use a budget share of taxable consumption of $B_\tau = 0.4$ based on ?. We set the elasticity of intertemporal substitution $\sigma = 0.5$ based on the meta study by ?. Anticipating our finding that relative spending on taxable and exempt goods is not much affected by the sales tax changes in the longer run, we choose a moderate intratemporal elasticity of substitution between taxable

and tax-exempt goods of $\epsilon = 0.4$. We use an income elasticity to sales tax changes of $\varepsilon_{\hat{y},\tau} = 0.1$ based on our own estimates. We assume that consumers learn about a sales tax increase three months in advance, $T = 3$, based on our analysis of news article coverage around sales tax changes. Finally, we use a monthly interest rate of $r = 0.08\%$ which corresponds to a 1% annual real risk-free rate. The only two parameters for which we do not have guidance from the literature are the measures of durability, δ_i . Targeting the response of taxable expenditures in the month of a sales tax increase, we find $\delta_\tau \approx 0.5$, which we also use to calibrate δ_e .

[Figure 5 about here]

Figure 5 graphs the model’s predicted response of consumption, inventories and expenditures to an anticipated 1% sales tax increase in month $T = 3$, both in log-levels and log-differences.

Discussion of model results TBA here.

The two main competing alternative explanations for the strong response of tax-exempt spending to sales tax changes besides shopping complementarity are (i) a high substitution elasticity ϵ between taxable and tax-exempt goods and (ii) confusion about which goods are exempt, i.e., non-salience of the exemption status. However, a high intratemporal substitution elasticity is inconsistent with the low long-run effect of relative spending on exempt and taxable goods. We discuss the salience of the exemption status in more detail below.

Identification of the household spending response

An important insight from the potential complementarity of taxable and tax-exempt expenditures via inventory “spill-overs” is that tax-exempt expenditures are not a valid control group in a household-level difference-in-difference research design. Instead, we exploit the high-frequency of the data around the date at which sales taxes change (i.e., at time T) using a before-after design and comparing treated households with households in untreated sales tax jurisdictions.

Conclusions

From 2004 to 2014, there were more than four thousand changes in state and local tax rates in the United States. Understanding if and how households respond differently to tax changes as compared to price changes has important implications both for tax

incidence but also more generally in structural models of household spending. This paper evaluates the impact of these tax changes on household retail spending across a number of dimensions.

Overall, we present evidence that households respond strongly to changes in sales taxes, both immediately following the tax change and also in the months leading up to the change. For instance, households bring spending forward to the months leading up to a tax increase and spend significantly less in the months afterward. This strong intertemporal substitution of spending is very short-lived despite the persistent change in the intertemporal price, suggesting that actual consumption behavior does not change significantly. Consistent with this hypothesis, we find the intertemporal substitution of spending is much larger for more storable or durable goods.

We also find that households are aware of other methods to avoid sales taxes, engaging in geographical arbitrage by increasing trips to locations with a lower sales tax rate after a tax increase in their home ZIP code and increasing the amount of purchases made from online merchants.

In the periods surrounding sales tax changes, we find that tax-exempt spending is affected to much the same degree as spending on taxable goods. To explain this seemingly irrational behavior, we build and calibrate a model of inventory and shopping complementarities where households rationally bundle purchases of different types of goods into single shopping trips. We show that this shopping complementarity mechanism has support in the data, with households possessing lower revealed shopping costs tend to bundle purchases less than households with higher costs.

While households, on average, respond strongly and rationally to changes in sales taxes, we also demonstrate that increases in the amount of information presented to households about upcoming changes to sales taxes tends to induce larger responses of household spending. These results imply that the salience of taxes is an important determinant of how households respond to taxes.

References


- Baker, Scott R. and Constantine Yannelis, “Income changes and consumption: Evidence from the 2013 federal government shutdown,” *Working Paper*, 2015.
- Baugh, Brian, Itzhak Ben-David, and Hoonsuk Park, “The ‘Amazon Tax’: Empirical Evidence from Amazon and Main Street Retailers,” *Working Paper*, 2015.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, “Salience and Consumer Choice,” *Journal of Political Economy*, 2013, 121 (5).
- Broda, Christian and David E. Weinstein, “Product Creation and Destruction: Evidence

- and Price Implications,” *American Economic Review*, 2010, 100 (3), 691–723.
- Cabral, Marika and Caroline Hoxby, “The Hated Property Tax: Salience, Tax Rates, and Tax Revolts,” *Stanford Working Paper*, 2011.
- Caplin, Andrew and Mark Dean, “Revealed Preference, Rational Inattention, and Costly Information Acquisition,” *American Economic Review*, 2015, 105 (7), 2183–2203.
- Cashin, David, “The Intertemporal Substitution and Income Effects of a Consumption Tax Rate Increase: Evidence from New Zealand,” *Working Paper*, 2014.
- , “The Household Expenditure Response to a Consumption Tax Increase,” *Working Paper*, 2015.
- Chen, Xiu, Harry M. Kaiser, and Bradley J. Rickard, “The Salience of Excise vs. Sales Taxes on Healthy Eating: An Experimental Study,” *Working Paper*, 2014.
- Chetty, Raj, Adam Looney, and Kory Kroft, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, 99 (4), 1145–1177.
- Clower, Robert W and Peter W Howitt, “The Transactions Theory of the Demand for Money: A Reconsideration,” *Journal of Political Economy*, 1978, pp. 449–466.
- D’Acunto, Francesco, Daniel Hoang, and Michael Weber, “The Effect of Unconventional Fiscal Policy on Consumption Expenditure,” *Working Paper*, 2016.
- Davis, Trevor, Dan Knoepfle, Stephen Sun, and Constantine Yannelis, “Greener on the Other Side: Estimating Consumer Sensitivity to Local Sales Tax Changes,” *Working Paper*, 2016.
- Einav, Liran, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan, “Sales Taxes and Internet Commerce,” *American Economic Review*, 2014, 104 (1), 1–26.
- , Ephraim Leibtag, and Aviv Nevo, “Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?,” *Quantitative Marketing and Economics*, 2010, (8), 207–239.
- Farhi, Emmanuel and Xavier Gabaix, “Optimal Taxation with Behavioral Agents,” *Working Paper*, 2015.
- Finkelstein, Amy, “EZTax: Tax Salience and Tax Rates,” *Quarterly Journal of Economics*, 2009.
- Gabaix, Xavier and David Laibson, “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 2006, 121 (2), 505–540.
- Gallagher, Kelly Sims and Erich Muehlegger, “Giving Green to Go Green? Incentives and Consumer Adoption of Hybrid Vehicle Technology,” *Journal of Environmental Economics and Management*, 2011, 61, 1–15.
- Goldin, Jacob, “Optimal Tax Salience,” *Working Paper*, 2014.
- Green, Daniel, Brian Melzer, Johnathan A. Parker, and Ryan Pfirmann-Powell, “Accel-

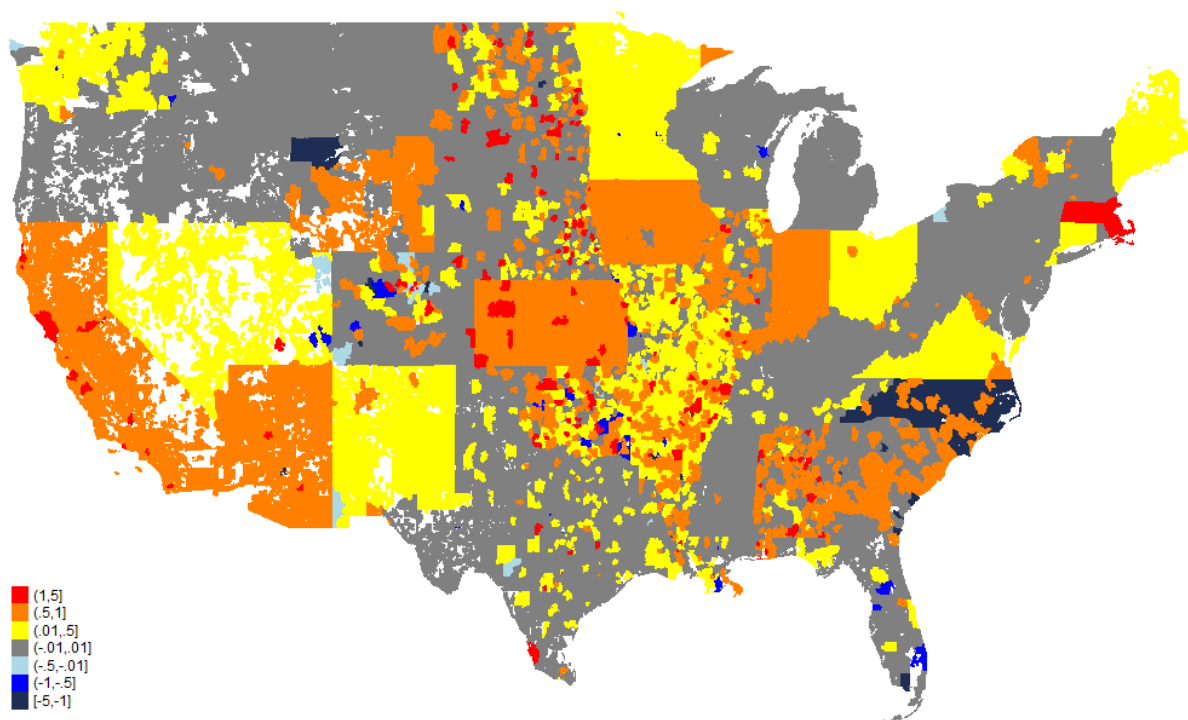
- erator or Brake? Microeconomic estimates of the ‘Cash for Clunkers’ and Aggregate Demand,” *Working Paper*, 2016.
- Havránek, Tomáš, “Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting,” *Journal of the European Economic Association*, 2015, *13* (6), 1180–1204.
- Hayashi, Andrew, “The Legal Salience of Taxation,” *81 U. Chi. L. Rev. Paper*, 2014, *1443*.
- Howitt, Peter, “Intertemporal Utility Maximization and the Timing of Transactions,” *American Economic Review*, 1977, *67* (2), 156–165.
- Kőszegi, Botond and Adam Szeidl, “A model of focusing in economic choice,” *The Quarterly Journal of Economics*, 2013, *128* (1), 53–104.
- Lobo, Miguel Sousa, Maryam Fazel, and Stephen Boyd, “Portfolio Optimization with Linear and Fixed Transaction Costs,” *Annals of Operations Research*, 2007, *152* (1), 341–365.
- Mian, Atif and Amir Sufi, “The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program,” *Quarterly Journal of Economics*, 2012, *127* (3), 1107–1142.
- Ott, Richard and David Andrus, “The Effects of Personal Property Taxes on Consumer Vehicle-Purchasing Decisions,” *Public Finance Review*, 2000, *28* (2), 134–152.
- Sims, Christopher A., “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, *50* (3), 665–690.
- Stephens-Davidowitz, Seth, “The Cost of Racial Animus on a Black Presidential Candidate: Evidence Using Google Search Data,” *Working Paper*, 2013.
- Taubinsky, Dmitry and Alexander R. Rees-Jones, “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment,” *Working Paper*, 2015.
- Taylor, Mac, “Why Have Sales Taxes Grown Slower Than the Economy?,” *Legislative Analyst’s Office Brief*, 2013.
- Woodford, Michael, “Inattentive Valuation and Reference-Dependent Choice,” *Working Paper*, 2012.

Figure 1 – ZIP Code Level Sales Tax Rates

(a) maximum sales tax rate levels, 2008-14

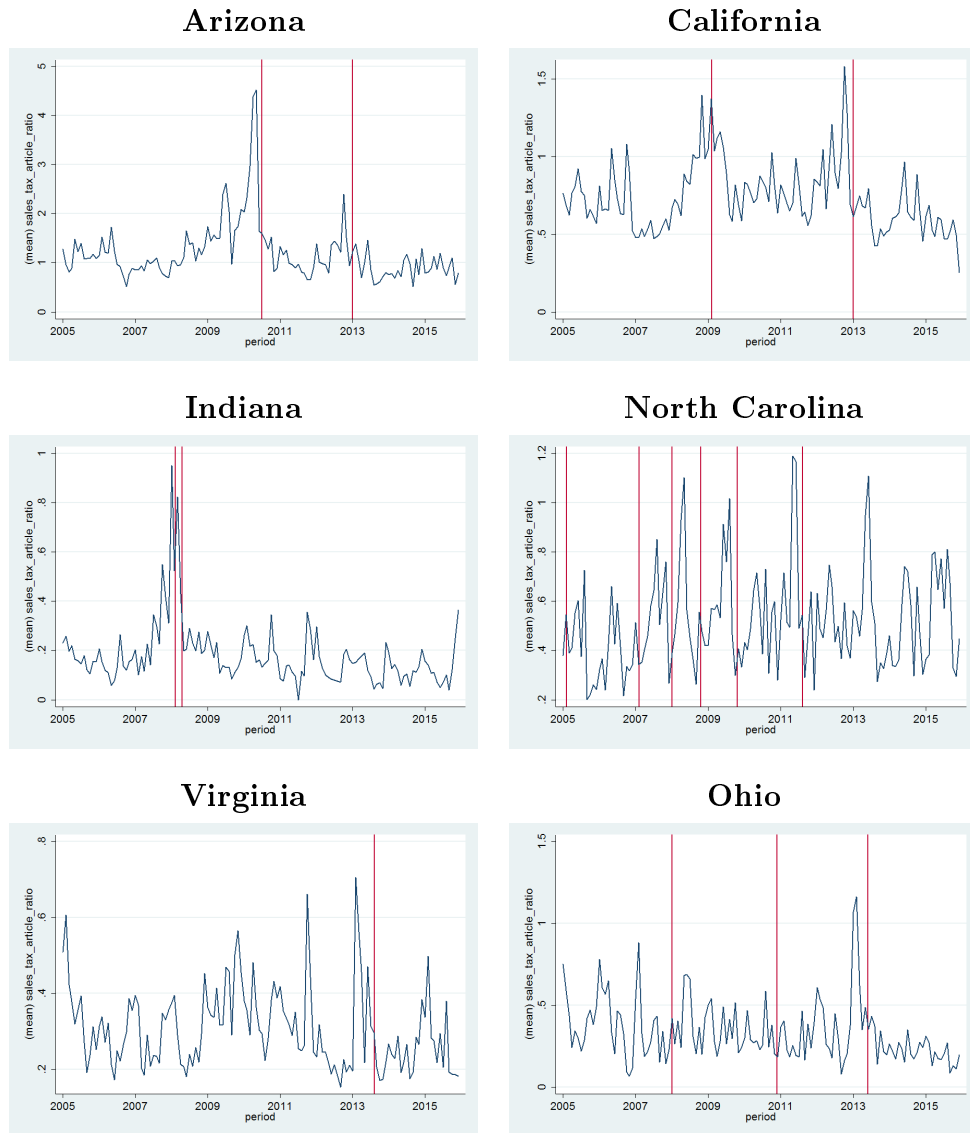


(b) maximum sales tax rate changes, 2008-14



Notes: Maps plot the maximum level (a) respectively change (b) of total sales tax rates in each five-digit ZIP code for years 2008-14, matching the sample period of the Nielsen Consumer Panel. Sales tax rates are expressed in percentages. Total sales tax rate changes may be driven by changes in state, city, county, or special district sales tax rates. White ZIP codes have missing sales tax rates or are not covered by Nielsen.

Figure 2 – Newspaper Coverage and Sales Tax Changes



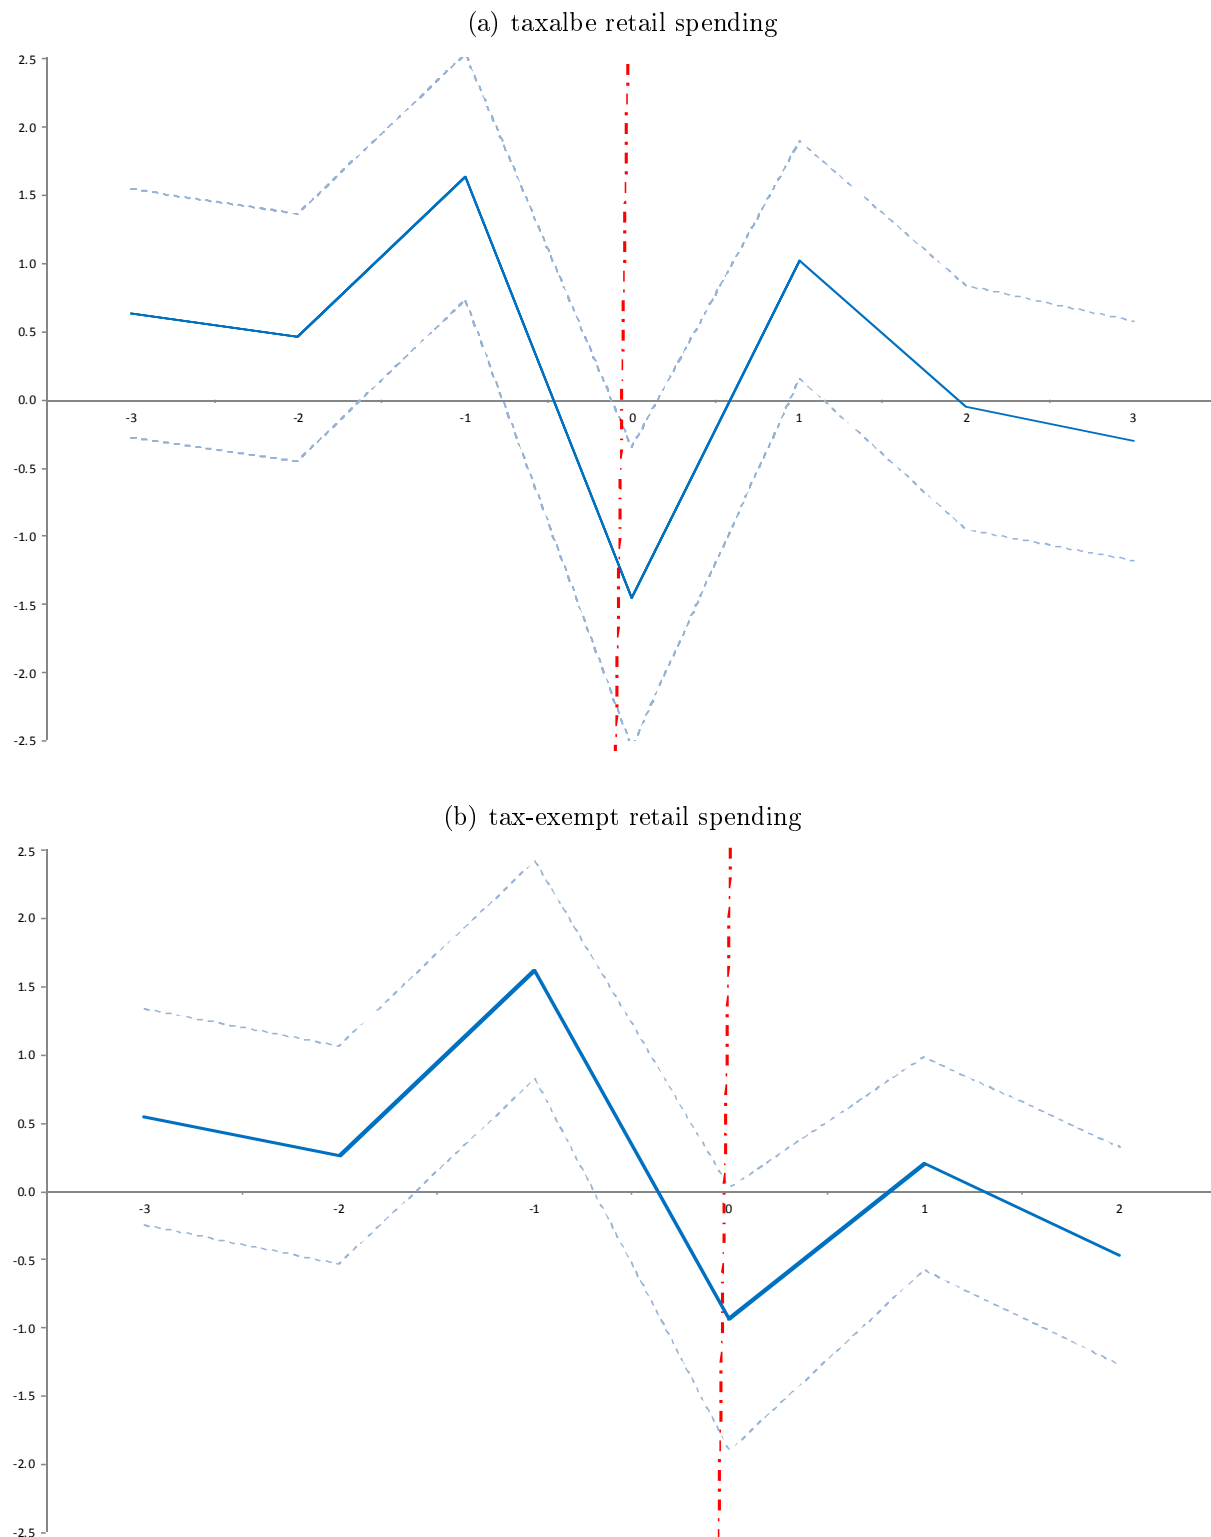
Notes: Figures plot news articles that contain the term ‘sales tax’ or ‘sales taxes’ as a fraction of all newspaper articles in a given month across newspapers in that state. Y-axis units are percentage points (eg. 0.2 = 0.2% of articles contain ‘sales tax’). Selected states shown. News articles taken from Access World News and cover approximately 3,000 US newspapers ranging from large national papers to local papers. Red vertical lines denote the dates of state-wide sales tax changes.

Figure 3 – Response of Newspaper Coverage and Google Searches to State Tax Increase



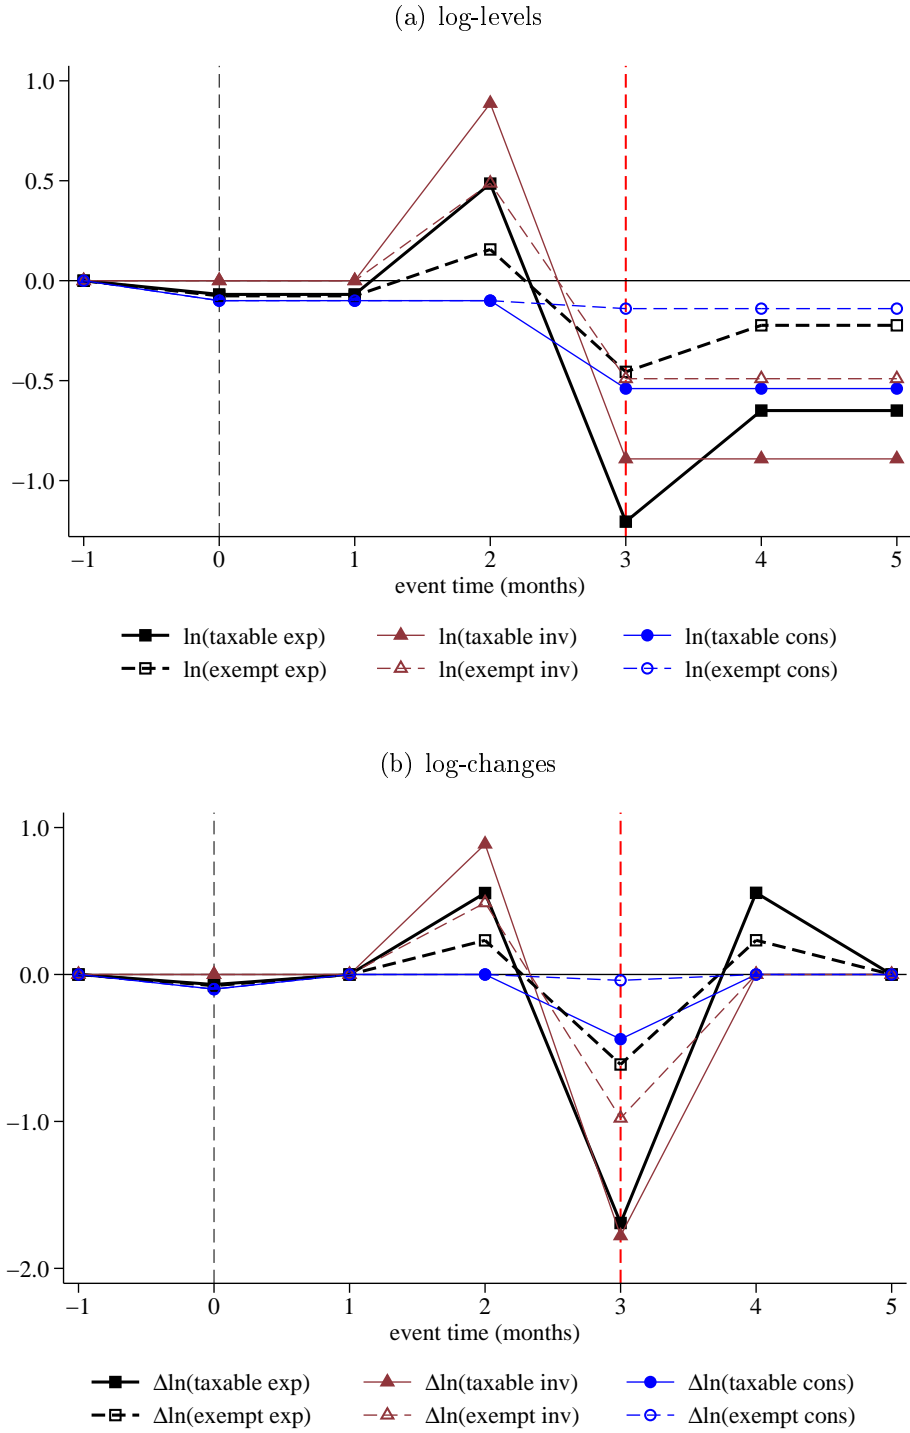
Notes: Top panel plots coefficients from a regression of the ratio of news articles that contain the term ‘sales tax’ or ‘sales taxes’ as a fraction of all newspaper articles in a given month across newspapers in that state. Y-axis units are percentage points. News articles taken from Access World News and cover approximately 3,000 US newspapers ranging from large national papers to local papers. Bottom panel plots coefficients from a regression of logged Google search activity from Google Trends. Y-axis units are percentage deviations from baseline. Household and period fixed effects are included. Standard errors clustered by state. Red vertical lines denote ‘time 0’, where a state level sales tax rate change occurs.

Figure 4 – Household-Level Retail Spending Response Around a State Sales Tax Change



Notes: Figure plots coefficients of a regression of the logged amount of pre-tax household retail spending on leads and lags of changes in state sales tax rates. Panel (a) shows the dynamic response of spending on taxable goods and (b) the response of tax-exempt spending to a sales tax rate increase. All coefficients are scaled to an increase in sales taxes of 1%. Dashed lines represent 95% confidence intervals from standard errors clustered at a state level. Regression period spans 2004-2014.

Figure 5 – Predicted response in model with baseline calibration



Notes: This figure shows the model's predicted response of consumption c , inventory s , and expenditures x to a 1% increase in the sales tax in period $T = 3$, which is announced in period $t = 0$. The baseline calibration uses $\sigma = 0.5$, $\epsilon = 0.4$, $\delta_\tau = \delta_e = 0.5$, $\theta_c/\theta_s = 1$, $\tilde{c}_\tau/\tilde{s}_\tau = 1.1$, $\tilde{c}_e/\tilde{s}_e = 1.6$, $B_\tau = 0.4$, $\varepsilon_{\hat{y},\tau} = 0.1$, $r = 0.08\%$.

Table 1. State and Local Changes in Sales Taxes by Year

	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	Total
State Sales Tax Changes	3	2	3	3	1	5	8	3	3	0	8	0	1	40
County Sales Tax Changes	-	-	-	-	-	260	208	68	59	65	92	88	115	955
City Sales Tax Changes	-	-	-	-	-	211	318	207	247	1,109	271	255	291	2,949
Other Local Sales Tax Changes	-	-	-	-	-	74	90	42	34	89	215	38	68	650
Total Sales Tax Changes	-	-	-	-	-	549	617	317	340	1,262	586	378	471	4,582

Notes: This table represents all state and local changes in sales tax rates in a given year and are mutually exclusive. A tax change and revision later the same year would be counted separately. In addition, expirations of sales tax changes are counted as a change, as well. ‘Other Local Sales Tax Changes’ include changes in administrative areas such as water districts, school districts, and other areas.

Table 2. Nielsen coverage statistics – 2011

Type of store	% of sales covered	
	% of Nielsen stores	for stores of this type
Convenience	6	2
Drug	35	55
Food	29	53
Liquor	1	1
Mass Merchandise	30	32

Notes: Data taken from Nielsen Homescan data documentation for 2011. Data used spans 2004-2014. Nielsen retailer coverage remained consistent across years though household sample size increased from under 40,000 to over 60,000.

Table 3: Response of Taxable Spending to a Sales Tax Increase

Dep. var.: $\Delta \ln(x)_{ht}$, monthly taxable retail spending	baseline	state tax only	Household characteristics	unemployment	drop NBER recession	state sales tax IV	state-period FE	all sales tax changes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Total sales tax rate	-1.916*** (0.601)		-1.914*** (0.601)	-1.959*** (0.601)	-1.821** (0.818)	-1.992*** (0.758)	-2.144*** (0.651)	-1.856*** (0.491)
Δ State sales tax rate		-2.051** (0.961)						
Δ Unemployment rate, state				-0.370** (0.180)	-0.526** (0.221)	-0.371** (0.180)	-0.991*** (0.225)	-0.370** (0.180)
Δ Unemployment rate, county				-0.108 (0.120)	-0.179 (0.144)	-0.108 (0.120)	-0.106 (0.120)	-0.116 (0.120)
- Household FE	YES	YES	YES	YES	YES	YES	YES	YES
- Year and Month FE	YES	YES	YES	YES	YES	YES	YES	YES
- Household characteristics	--	--	YES	YES	YES	YES	YES	YES
- State-year FE	--	--	--	--	--	--	YES	--
- State-month FE	--	--	--	--	--	--	YES	--
Observations	4,137,927	5,928,468	4,137,927	4,137,927	3,285,747	4,137,927	4,137,927	4,151,336
R-squared	0.013	0.012	0.013	0.013	0.014	0.013	0.013	0.013

Notes: Total sales tax rates combine all sales taxes within a ZIP code, including state, county, city, and special districts. The dependent variable is monthly changes in logged household taxable spending as measured by Nielsen Consumer Panel data. Taxability of household spending is defined at a state level depending on what categories of goods are exempt from sales taxes (e.g., groceries, clothing, medication). For robustness, the dependent variable is winsorized at the 1% level. Other household characteristics include fixed effects for income bins and family size. Regressions span 2004-2014 for state sales tax rate changes (column 2) and 2008-2014 for total sales tax rate changes. Column 5 excludes the Great Recession based on NBER recession dates, which is January 2008 to June 2009. Column 6 uses state sales tax rate changes as an instrument to deal with potential measurement error in local sales tax rate changes. Robust standard errors in parentheses adjust for arbitrary within-household correlations and heteroskedasticity and are clustered at the ZIP code for total sales tax rate changes and at the state level for state sales tax rate changes.

Table 4: Tax Avoidance Strategies - Jurisdictional Tax Arbitrage and Online Spending

<i>Dependent variables:</i>	<i>Δ fraction spent in alternative ZIP-3 code</i>		<i>Δ ln(online spending)</i>
	(1)	(2)	(3)
Δ Total sales tax rate	0.015 (0.077)	-0.169*** (0.057)	
Δ Total sales tax rate x (avg. frac spent in alt ZIP-3		2.756*** (1.055)	
Δ State sales tax rate			1.605** (0.763)
- Household FE	YES	YES	YES
- Year and Month FE	YES	YES	YES
Observations	4,231,049	4,231,049	6,871,487
R-squared	0.000	0.000	0.004

Notes: Total sales tax rates combine all sales taxes within a ZIP code, including state, county, city, and special districts. The dependent variable in columns 1 and 2 is the change in the monthly fraction of a household's retail spending in an alternative tax jurisdiction outside the household's residential 3-digit ZIP code. The dependent variable in column 3 is the change in the log of total online spending, including mail orders. For robustness, dependent variables are winsorized at the 1% level. Regressions span 2004-2014 for state sales tax rate changes and 2008-2014 for total sales tax rate changes. Columns 1 and 2 use state tax rate changes as an instrument to deal with potential measurement error in local sales tax rate changes. Robust standard errors in parentheses adjust for arbitrary within-household correlations and heteroskedasticity and are clustered at the ZIP code or state level.

Table 5: Spending Response by Product Storability

Dep. var.: $\Delta \ln(x)_{ht}$, monthly retail spending by product group and state	baseline	by inverse shopping freq.	durability indicator only	<i>anticipation and intertemporal substitution</i>			<i>top and bottom product groups by purchase freq.</i>	
				baseline	by inverse shopping freq.	by durability only	product group	avg # of purch. per month
	(1)	(2)	(3)	(4)	(5)	(6)		
Δ State sales tax rate	-2.271*** (0.654)	-2.070*** (0.580)	-1.925*** (0.660)	-2.271*** (0.654)	-2.068*** (0.579)	-1.922*** (0.661)	BREAD AND BAKED GOODS	2.831
Δ State sales tax rate, lead				0.263 (0.942)	0.012 (0.941)	-0.019 (1.007)	FRESH PRODUCE	2.508
Δ State sales tax rate x $\ln(\text{shopping cycle})$		-1.665* (0.930)			-1.690* (0.943)		MILK	2.380
Δ State sales tax rate x $\ln(\text{shopping cycle})$, lead					1.996** (0.744)		CARBONATED BEVERAGES	1.942
Δ State sales tax rate x $I(\text{durable})$			-2.153 (2.626)			-2.174 (2.647)
Δ State sales tax rate x $I(\text{durable})$, lead						1.746* (0.945)	FLOUR	0.209
- Product FE	YES	YES	YES	YES	YES	YES	FRAGRANCES - WOMEN	0.081
- State FE	YES	YES	YES	YES	YES	YES	PHOTOGRAPHIC SUPPLIES	0.076
- Year and Month FE	YES	YES	YES	YES	YES	YES		
Observations	493,649	493,649	493,649	493,649	493,649	493,649	MEN'S TOILETRIES	0.075
Sample mean of interaction variable		0.101	0.158		0.101	0.158	Sample mean	1.140
R-squared	0.078	0.078	0.078	0.078	0.078	0.078	Sample standard deviation	0.714

Notes: The dependent variable is the monthly log change in total retail spending by product group and state across all households in the Nielsen Consumer Panel, 2004-2014. We drop product groups which on average have less than \$500 total sales per state. We also drop "magnet data", leaving 108 unique product groups. All regressions are estimated using least squares weighted by average sales per product group. Robust standard errors in parentheses adjust for arbitrary within-state correlations and heteroskedasticity and are clustered at the state level.

Table 6: Response of Tax-Exempt Spending to a Sales Tax Increase

Dep. var.: $\Delta \ln(x)_{ht}$, monthly tax-exempt retail spending	baseline	state tax only	Household characteristics	unemployment	drop NBER recession	state sales tax IV	state-period FE	all sales tax changes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Total sales tax rate	-1.280*** (0.476)		-1.279*** (0.476)	-1.221** (0.476)	-1.769*** (0.649)	-1.612*** (0.595)	-1.100** (0.518)	-1.038*** (0.396)
Δ State sales tax rate		-1.504** (0.614)						
Δ Unemployment rate, state				0.742*** (0.145)	0.848*** (0.179)	0.740*** (0.145)	0.852*** (0.182)	0.747*** (0.145)
Δ Unemployment rate, county				-0.105 (0.099)	-0.095 (0.120)	-0.105 (0.099)	-0.141 (0.098)	-0.105 (0.099)
- Household FE	YES	YES	YES	YES	YES	YES	YES	YES
- Year and Month FE	YES	YES	YES	YES	YES	YES	YES	YES
- Household characteristics	--	--	YES	YES	YES	YES	YES	YES
- State-year FE	--	--	--	--	--	--	YES	--
- State-month FE	--	--	--	--	--	--	YES	--
Observations	4,095,538	5,865,352	4,095,538	4,095,538	3,250,600	4,095,538	4,095,538	4,108,809
R-squared	0.014	0.012	0.014	0.014	0.015	0.014	0.015	0.014

Notes: Total sales tax rates combine all sales taxes within a zip-code, including state, county, city, and special districts. Dependent variable is monthly changes in logged household tax-exempt spending as measured by Nielsen Consumer Panel data. Tax-exemption status of household spending is defined at a state level. For robustness, the dependent variable is winsorized at the 1% level. Other household characteristics include fixed effects for income and family size. Regressions span 2004-2014 for state sales tax rate changes (column 2) and 2008-2014 for total sales tax rate changes. Column 5 excludes the Great Recession based on NBER recession dates, which is January 2008 to June 2009. Column 6 uses state sales tax rate changes as an instrument to deal with potential measurement error in local sales tax rate changes. Robust standard errors in parentheses adjust for arbitrary within-household correlations and heteroskedasticity and are clustered at the ZIP code for total sales tax rate changes and at the state level for state sales tax rate changes.

Table 7: Evidence of Shopping Complementarities

<i>Dependent variable:</i>	response of shopping trips		revealed cost approach: heterogeneity in shopping frequency				"placebo" test: online shopping	
	$\Delta \ln(\# \text{ of trips})$	$\Delta \ln(\# \text{ of store visits per day})$	<i>frequent shoppers (low cost)</i>		<i>infrequent shoppers (high cost)</i>		$\Delta \ln(\text{taxable})$	$\Delta \ln(\text{exempt})$
			$\Delta \ln(\text{exempt})$	$\Delta \ln(\text{taxable})$	$\Delta \ln(\text{exempt})$	$\Delta \ln(\text{taxable})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Total sales tax rate	-1.805*** (0.435)	-0.766** (0.305)	-0.001 (0.703)	-2.048** (0.845)	-2.076* (1.097)	-2.358* (1.343)		
Δ State sales tax rate							1.492** (0.686)	0.247 (0.508)
- Household FE	YES	YES	YES	YES	YES	YES	YES	YES
- Year and Month FE	YES	YES	YES	YES	YES	YES	YES	YES
Observations	4,231,067	4,231,067	1,086,935	1,091,679	934,665	951,895	6,871,997	6,879,154
R-squared	0.017	0.019	0.016	0.016	0.020	0.017	0.004	0.003

Notes: Total sales tax rates combine all sales taxes within a ZIP code, including state, county, city, and special districts. The dependent variable is the monthly log-change in the number of shopping trips (column 1) resp. stores visited per day (column 2). Frequent shoppers in columns 3 and 4 are consumers with average monthly trips above the 75th percentile (19 trips), while infrequent shoppers in columns 5 and 6 have average monthly trips below the 25th percentile (9 trips). For robustness, the dependent variables are winsorized at the 1% level. Regressions span 2008-2014 for total sales tax rate changes and 2004-2014 for state sales tax rate changes. Robust standard errors in parentheses adjust for arbitrary within-household correlations and heteroskedasticity and are clustered at the ZIP code or state level.

Table 8: Saliency and Announcement Effects

Dependent variable:	saliency effects						announcement effects using month of ballot proposition	
	newspaper coverage			ballot-induced tax changes			$\Delta \ln(\text{total})$	$\Delta \ln(\text{total})$
	$\Delta \ln(\text{total})$	$\Delta \ln(\text{taxable})$	$\Delta \ln(\text{exempt})$	$\Delta \ln(\text{total})$	$\Delta \ln(\text{taxable})$	$\Delta \ln(\text{exempt})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ State sales tax rate	-1.610*** (0.561)	-2.071** (0.973)	-1.460** (0.565)	-1.541** (0.647)	-2.107* (1.101)	-1.277** (0.608)		
Δ State sales tax rate x newspaper coverage	-77.025** (34.182)	-57.484 (70.606)	-114.317*** (42.371)					
Newspaper coverage of state sales tax changes	-0.143 (0.184)	-0.372* (0.212)	0.241 (0.228)					
Δ State sales tax rate x I(ballot proposition)				-5.226*** (0.986)	-4.515** (1.916)	-6.324*** (0.831)		
I(state tax ballot proposition)				0.027*** (0.006)	0.030*** (0.009)	0.025*** (0.005)		
State sales tax rate change proposed							-0.605* (0.351)	-2.039 (1.604)
State sales tax rate change proposed x I(ballot prop. failed)								1.687 (1.642)
I(ballot prop. failed)								-0.000 (0.007)
- Household FE	YES	YES	YES	YES	YES	YES	YES	YES
- Year and Month FE	YES	YES	YES	YES	YES	YES	YES	YES
Observations	5,928,468	5,928,468	5,865,352	5,928,468	5,928,468	5,865,352	5,941,471	5,941,471
R-squared	0.013	0.012	0.012	0.013	0.012	0.012	0.013	0.013

Notes: The dependent variables are monthly changes in logged household spending as measured by Nielsen Consumer Panel data. Taxability and tax-exemption status of household spending is defined at a state level depending on what categories of goods are exempt from sales taxes (e.g., groceries, clothing, medication). Columns 1-3 interact changes in sales tax rates with the level of newspaper coverage (measured as the ratio of articles mentioning sales taxes to the total number of articles in newspapers within the state covered by Access World News). Columns 4-6 interact changes in sales tax rates with an indicator for whether the change in state sales tax rates was driven by a ballot measure (as opposed to being enacted by the legislature). Columns 7 and 8 use, as independent variables, indicators for dates when ballot initiatives that would change state sales tax rates were voted on (as opposed to the dates they were enacted). Column 8 interacts these indicators with another indicator that signifies the ballot not being successfully passed (and thus resulting in no change in sales tax rates). For robustness, the dependent variables are winsorized at the 1% level. Regressions span years 2004-2014. Robust standard errors in parentheses adjust for arbitrary within-household correlations and heteroskedasticity and are clustered at the state level.

Table 9: Robustness - Quantity and Price Response to a Sales Tax Increase

<i>dependent variable:</i>	quantities				prices			
	$\Delta \ln(\text{taxable})$	$\Delta \ln(\text{exempt})$	$\Delta \ln(\text{taxable})$	$\Delta \ln(\text{exempt})$	$\Delta \ln(\text{retail price})$	$\Delta \ln(\text{wholesale price})$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Total sales tax rate	-2.330*** (0.479)	-1.458*** (0.458)			-0.215*** (0.036)		-0.008* (0.004)	
Δ State sales tax rate			-2.245** (0.908)	-1.744*** (0.566)		-0.171** (0.069)		-0.007 (0.015)
- Household FE	YES	YES	YES	YES				
- Year and Month FE	YES	YES	YES	YES	YES	YES	YES	YES
- Product FE					YES	YES	YES	YES
- ZIP-3 FE					YES	YES	YES	YES
Observations	4,140,969	4,142,698	5,928,529	5,928,499	4,333,000	5,862,621	4,333,000	5,862,621
R-squared	0.014	0.014	0.013	0.013	0.011	0.010	0.189	0.177

Notes: Total sales tax rates combine all sales taxes within a ZIP code, including state, county, city, and special districts. Dependent variables in columns 1 to 4 are monthly changes in logged quantities (items) purchased by each household in the Nielsen Consumer Panel. Dependent variables in column 5 to 8 are monthly changes in sales-weighted average prices by product group and ZIP-3 code for all retailers in the Nielsen Retail Scanner Panel. For robustness, the dependent variables are winsorized at the 1% level. Regressions span 2004-2014 for state sales tax rate changes using the Nielsen Consumer Panel (columns 3 and 4) respectively 2006-2014 using the Nielsen Retail Scanner Panel (columns 6 and 8), and 2008-2014 for total sales tax rate changes (columns 1, 2, 5, and 7). Robust standard errors in parentheses adjust for arbitrary within-product correlations and heteroskedasticity and are clustered at the ZIP-3 code for total sales tax rate changes and at the state level for state sales tax rate changes.

Online Appendix of “Shopping for Lower Sales Tax Rates”

Scott R. Baker Lorenz Kueng

Pricing Data

Nielsen Retail Scanner (Retailer Data)

With the Nielsen Retail Scanner (NRS) data, price and quantity information is available at the store level for each UPC carried by a covered retailer and span the years 2006-2014. Nielsen provides the location of the stores at the three-digit zip code level (eg. 602 instead of 60208), county and Nielsen’s Designated Market Area(DMA). An average (quantity weighted) price is reported, by UPC, for each store every week.¹⁷

In the NRS data, units are consistently standardized and most products are measured in ounces (OZ, 51%), count (CT, 45%) or ml (ML, 2%). NRS covers 125 product groups with more than 3.2 million individual UPCs. Aside from grocery items, the dataset also includes cosmetics, alcohol and general merchandise (homewares, office supplies, some electronics, garden supplies). The types of stores are detailed in Table 2.

PromoData (Wholesale Prices)

We use PromoData to measure wholesale prices for grocery and retail goods. Promo obtains its information from one (confidential) major wholesaler in each market.¹⁸ One downside to this approach is that, since no single wholesaler carries every SKU in a given market, information about the universe of goods is not observed. Overall, Promo prices are available for 32 markets after removing redundant markets and combining overlapping markets.¹⁹

Data on wholesale prices are available from 2006 - 2012. However, during 2012 the

¹⁷For a given store, coverage over time is stable and relatively complete across all years. Unit prices are calculated as $price/(prmult \times size1_amt)$; see Note E on p.15 of ‘Retail Scanner Dataset Manual 05.25.2015.pdf.’

¹⁸By only using one wholesaler Promo relies on the Robinson–Patman Anti-Price Discrimination Act of 1936 that prohibits price discrimination. In particular, it prevents wholesalers from offering special discounts to large chain stores but not to other, smaller retailers.

¹⁹Leveraging this regional information provides additional variation but introduces more measurement error given less complete coverage in any given market both with respect to corresponding Nielsen product groups in the cross-section and time-series coverage of specific products.

data loses a significant amount of coverage. For this reason, we perform robustness tests excluding 2012 data from our sample. PromoData contains all changes in price or deals that are run by the wholesaler. Thus, we take prices as constant between observations, based on the last observed price data. We then are able to collapse prices to a monthly level for each product group. To arrive at consistent unit prices within type of product (eg. product groups), we scale the observed wholesale prices by the number of goods in a ‘pack’ and by the size of the unit (eg. number of ounces in a candy bar and number of candy bars in a box). To make meaningful unit price comparisons we need to know the units associated with each good. Unfortunately unit information is often not provided or is inconsistently coded (e.g. CT, PACK, EACH, OZ, O etc.). We use the modal unit within UPC to impute missing values. The intuition is that if a product is recorded as being measured in OZ most of the time units are reported, it is probably measured in OZ.

Matching Wholesaler and Retailer Data

Given the large number of products in the retailer dataset we aggregate retail unit prices to the product group level before matching with wholesale prices. We assign products in the wholesaler data to Nielsen product groups by matching at the UPC level. The mapping is not one-to-one due to differences in end-digits when shifting to UPCs of different levels of granularity (eg. some are reported with retailer specific end-digits, etc.). This leads to multiple Nielsen UPCs corresponding to a single Promo UPC for some goods. However, this appears to have little effect when merging Nielsen product groups to their Promo equivalents.

As a consistency check we also match retail and wholesale prices by UPC at a single point in time. The implied markup distribution supports the accuracy of both the raw data and our unit price calculations, with 90% of markups falling between -7% and 135%. We calculate Promo coverage of Nielsen product groups as the percentage of UPCs in each Nielsen product group that can be found in Promo. Overall, we see that about 4% of overall UPCs in Nielsen are also covered directly in the wholesale data for a given market. Aggregating across markets to the national level, this coverage increases somewhat.

The two datasets are merged based on the weekly date. That is, Promo prices are those associated with the week containing the Nielsen week-ending Saturday. For a Nielsen retailer using a 7-day period ending on Saturday the periods correspond closely. However, as mentioned above this is not the case for all retailers. For a retailer using a Thursday to Wednesday week, the Nielsen prices would pre-date the Promo prices by a few days.

Comparing unit prices is not completely straightforward as Promo units are missing for many products. As discussed above, we impute some missing units based on the modal

unit reported in Promo for that UPC. When merging, we retain only UPCs for which the imputed Promo unit matches the Nielsen unit. A coarse attempt is made to standardize the more common Promo units before matching. In particular we assume O and Z refers to OZ and $C, CNT, PK, EA, EACH, STK, ROL, RL, PC, \#, CTN$ refer to CT .

Details about the Model

Augmented Baumol-Tobin model with trip complementarity

Here we develop a possible micro-foundation of the transaction cost function θ with inventory complementarity. We augment a classic Allais-Baumol-Tobin model that trades off reductions in shopping trips against holding costs of capital tied up in inventories at below market interest rate in steady state. We assume that this model is solved in steady state, i.e., consumers take their optimal consumption levels \tilde{c}_i as given,²⁰ and that individuals consume each good at constant rate within each period, which is consistent with the intertemporal Euler equation when $\rho = r$. Costs from shopping trips are linear in the number of trips per period, $\theta_s \times f(s)$, where θ_s are the opportunity cost in dollars of a single trip and f is the frequency of trips which depends inversely on the level of inventory, $f' < 0$, and which we assume to be convex, $f'' > 0$. Abstracting from integer constraints, we assume that a consumer pays this cost in each period in proportion to the shopping frequency. We can calculate the inventory cycle time T in the (continuous-time) Baumol-Tobin model with inventory depreciating at rate $\lambda_i \equiv 1 - \delta_i$ and constant demand c_i by solving the ordinary differential equation $ds_i(t) = [-\lambda_i s_i(t) - \tilde{c}_i]dt$, with boundary conditions $s_i(0) = s_i$ and $s_i(T) = 0$. The frequency of purchases per period, the inverse of the cycle time, is

$$f_i(s_i|\tilde{c}_i, \lambda_i) = \frac{\lambda_i}{\ln\left(1 + \lambda_i \frac{s_i}{\tilde{c}_i}\right)} \xrightarrow{\lambda_i \rightarrow 0} \frac{\tilde{c}_i}{s_i}.$$

The frequency has the properties $\frac{\partial f_i}{\partial s_i} < 0$, $\frac{\partial^2 f_i}{\partial s_i^2} > 0$, $\frac{\partial f_i}{\partial \tilde{c}_i} > 0$, and $\frac{\partial^2 f_i}{\partial \tilde{c}_i^2} \xrightarrow{\lambda_i \rightarrow 0} 0$.

We augment these costs to take into account that the consumer can save on shopping trips when temporarily deviating from steady state inventory of one good by also deviating in the same direction for the other good. For instance, when bringing spending forward to periods before a tax or price increase, the consumer can also save on future trips for tax-exempt goods by stocking up on those goods in the same period. Hence, for calibration, these costs are of the same order of magnitude as the opportunity cost of

²⁰ This seems sensible and it maintains the separability of the intertemporal and intratemporal allocation problems.

trips. We assume that the saving on trips costs for tax-exempt goods is proportional to the additional cost. For instance, if the consumer's shopping cost increases by $k\%$ due to temporarily increasing the inventory of taxable goods by 1% above steady state, we assume that she gets an offsetting credit of $k\%$ when also stocking up the inventory of exempt goods by 1% . Hence, this assumption avoids the double counting of shopping transactions costs when synchronizing additional out-of-steady-state purchases of exempt and taxable goods. For tractability we take a first-order approximation of the shopping-trip frequency with $\delta = 1$ around the steady state,

$$\frac{\tilde{c}_i}{s_{it}} - \frac{\tilde{c}_i}{\tilde{s}_i} \approx -\frac{\tilde{c}_i}{\tilde{s}_i} \left(\frac{s_{it} - \tilde{s}_i}{\tilde{s}_i} \right).$$

Putting these two sources of inventory costs (or rather benefits) together results in the inventory cost function (19).

Announcement effect on consumption

The term D in equation (12) to announcement of a permanent sales tax change in period T responds as

$$\begin{aligned} -\frac{d \ln D}{d \ln(1 + \tau)} &= \frac{1}{D} \left\{ P_0^\sigma \sum_{t=0}^{\infty} e^{-rt} \frac{\partial P_t^{1-\sigma}}{\partial \ln(1 + \tau_t)} + \mathbb{1}_{T=0} \frac{\partial P_0^\sigma}{\partial \ln(1 + \tau_0)} \sum_{t=0}^{\infty} e^{-rt} P_t^{1-\sigma} \right\} \\ &= \frac{1}{D} \left\{ (1 - \sigma) B_\tau P_0^\sigma \sum_{t=T}^{\infty} e^{-rt} P_t^{1-\sigma} + \mathbb{1}_{T=0} \sigma B_\tau P_0^\sigma D \right\} \\ &= \frac{1}{D} \left\{ (1 - \sigma) B_\tau e^{-rT} \underbrace{\left(\frac{P_0}{P_T} \right)^\sigma}_{C_T/C_0=1} \underbrace{P_T^\sigma \sum_{t=T}^{\infty} e^{-r(t-T)} P_t^{1-\sigma}}_D + \mathbb{1}_{T=0} \sigma B_\tau P_0^\sigma D \right\} \\ &= -e^{-rT} B_\tau (1 - \sigma) - \sigma B_\tau \mathbb{1}_{T=0}. \end{aligned}$$

Define the elasticity of income net of inventory costs as

$$\begin{aligned} \varepsilon_{\hat{y}, \tau} \equiv \frac{d \ln \hat{y}_t}{d \ln(1 + \tau)} &= \frac{y}{\hat{y}} \frac{d \ln y}{d \ln(1 + \tau)} + \frac{1}{\hat{y}} \sum_i \left[\left(\theta_h p_i \tilde{s}_i - \theta_s \frac{\tilde{c}_i}{\tilde{s}_i} + p_i \tilde{s}_i (1 - \delta_i) \right) \frac{d \ln \tilde{s}_i}{d \ln(1 + \tau)} \right. \\ &\quad \left. + (1 - \delta_i) p_i \tilde{s}_i \frac{d \ln p_i}{d \ln(1 + \tau)} + \theta_s \frac{\tilde{c}_i}{\tilde{s}_i} \frac{d \ln \tilde{c}_i}{d \ln(1 + \tau)} \right] \mathbb{1}_{t \geq T}, \end{aligned}$$

where $\frac{d \ln \tilde{z}_i}{d \ln(1 + \tau)}$ denotes the change in z_i from the old to the new steady state. This expression is constant since we are evaluating it at the steady state, i.e., it reflects the relative change in net income from the old to the new steady state. Using this definition,

we can express the response of wealth to a sales tax change as

$$\frac{d \ln W}{d \ln(1 + \tau)} = \frac{1}{W} \sum_{t=0}^{\infty} e^{-rt} \hat{y}_t \frac{d \ln \hat{y}}{d \ln(1 + \tau)} \approx e^{-rT} \varepsilon_{\hat{y}, \tau} \frac{1}{W} \overbrace{\sum_{t=T}^{\infty} e^{-r(t-T)} y_t}^W = e^{-rT} \varepsilon_{\hat{y}, \tau}.$$

The approximation is due to the fact that the definition of the elasticity ignores two effects that already happen before the taxes change at date T , which are given by

$$\frac{1}{W} \left\{ \sum_i \left[p_i s_i \left(\delta_i \frac{d \ln s_{iT}}{d \ln(1 + \tau)} + (1 - \delta_i) \frac{d \ln s_{i,T-1}}{d \ln(1 + \tau)} \right) \right] + \sum_{t=0}^{T-1} e^{-rt} \sum_i \theta_s \frac{\tilde{c}_i}{\tilde{s}_i} \frac{d \ln c_{it}}{d \ln(1 + \tau)} \right\} \approx 0.$$

Both of these effects are approximately zero as a fraction of total wealth W .²¹ The first takes into account that inventories first increase in period $T - 1$ before permanently decreasing to the new steady state level when taxes change in period T . The second accounts for the fact that consumption already changes between period 0 and $T - 1$, and the level in periods 0 to $T - 1$ is different from both initial and final steady state due to intertemporal substitution. This change in consumption—which we assume the consumer ignores for tractability—in turn affects the marginal shopping cost through a change in the shopping trip frequency, $\frac{c_{it}}{s_{it}}$. Instead of the actual shopping frequency, the consumer compares inventories relative to the (new) steady state consumption \tilde{c}_i , i.e., $\theta_h p_{it} s_{it} + \theta_s \frac{\tilde{c}_i}{s_{it}}$, while the “true” cost uses the actual frequency and is thus $\theta_h p_{it} s_{it} + \theta_s \frac{\tilde{c}_i}{s_{it}} + \theta_s \frac{c_{it} - \tilde{c}_i}{s_{it}}$. Hence, the third term captures the effect of the change in consumption on the inventory cost function in the transition from period 0 to $T - 1$, $\frac{d}{d \ln(1 + \tau)} \left[\theta_s \frac{c_{it} - \tilde{c}_i}{s_{it}} \right] = \theta_s \frac{\tilde{c}_i}{\tilde{s}_i} \frac{d \ln c_{it}}{d \ln(1 + \tau)}$. Combining $\frac{d \ln W}{d \ln(1 + \tau)} - \frac{d \ln D}{d \ln(1 + \tau)}$ yields Slutsky equation (14).

Log-linearizing expenditure changes

$$\begin{aligned} \Delta \ln x_{it} &\approx \frac{\Delta x_{it}}{x_{i,t-1}} = \frac{s_{i,t-1}}{x_{i,t-1}} \left[\frac{\Delta s_{it}}{s_{i,t-1}} - \frac{s_{i,t-1}}{s_{i,t-2}} \delta_i \frac{\Delta s_{i,t-1}}{s_{i,t-2}} + \frac{c_{i,t-1}}{s_{i,t-1}} \frac{\Delta c_{it}}{c_{i,t-1}} \right] \\ &= \frac{s_{i,t-1}}{x_{i,t-1}} \left[\Delta \ln s_{it} - \frac{s_{i,t-1}}{s_{i,t-2}} \delta_i \Delta \ln s_{i,t-1} + \frac{c_{i,t-1}}{s_{i,t-1}} \Delta \ln c_{it} \right] \\ &\approx \frac{\tilde{s}_i}{\tilde{x}_i} \left[\Delta \ln s_{it} - \delta_i \Delta \ln s_{i,t-1} + \frac{\tilde{c}_i}{\tilde{s}_i} \Delta \ln c_{it} \right], \end{aligned}$$

which is equation (22). In steady state with $s_{it} = \tilde{s}_i$ and $c_{it} = \tilde{c}_i$, equation (6) implies that $\tilde{x}_i = (1 - \delta_i) \tilde{s}_i + \tilde{c}_i$, hence the steady-state inventory-to-expenditure ratio is $\frac{\tilde{s}_i}{\tilde{x}_i} =$

²¹ Since these additional wealth effects are all capitalized already at period 0, they do not affect the response in the period before and after an anticipated sales tax change. Instead, they affect the interpretation and calibration of the wealth effect on C_0 . Hence, adding both terms instead of ignoring them leaves our results unchanged but would complicate the formula in (14).

$$[1 - \delta_i + \frac{\tilde{c}_i}{\tilde{s}_i}]^{-1}.$$

Allowing for durable goods in the storable goods framework

Durables depreciate for two main reasons: either because of utilization (e.g., driving more miles with a car per period) or because of the passing of time (e.g., corrosion or obsolescence). Denote ζ the rate of depreciation due to the passing of time and $\gamma_{d,t}$ the depreciation due to utilization of the stock of durables in period t . Hence, of the initial stock of durables d_{t-1} available in at the beginning of period t , only $1 - \zeta - \gamma_{d,t}$ is available in the next period. $x_{d,t}$ are expenditures on durables (maintenance or new acquisitions). Defining the consumption of durables in period t as $c_{d,t} = \gamma_{d,t}d_{t-1}$ and $\delta_d = 1 - \zeta$, we obtain the same law of motion for the stock of durables as for storables,

$$\begin{aligned} d_t &= (1 - \zeta - \gamma_{d,t})d_{t-1} + x_{d,t} \\ &= \delta_d d_{t-1} - c_{d,t} + x_{d,t} . \end{aligned}$$