NBER WORKING PAPER SERIES

SHOPPING FOR LOWER SALES TAX RATES

Scott R. Baker
Stephanie Johnson
Lorenz Kueng

Working Paper 23665
http://www.nber.org/papers/w23665

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2017

An early version from June 2016 circulated with the title "Margins of Adjustment to Price and Sales Tax Changes: New Evidence from Retailer-Customer Linked Panel Data." We would like to thank our discussants Jonathan Parker and Inessa Liskovich, and participants at seminars at the NBER Summer Institute, Arizona State, Berkeley, Bern, Columbia, Frankfurt, Minneapolis Fed, Minnesota, Munich, Northwestern, NYU, Yale and Venice for their comments. This research received financial support from the Alfred P. Sloan Foundation through the NBER Household Finance small grant program and from the Financial Institutions and Markets Research Center at the Kellogg School of Management. 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Shopping for Lower Sales Tax Rates
Scott R. Baker, Stephanie Johnson, and Lorenz Kueng
NBER Working Paper No. 23665
August 2017
JEL No. D12,E21,H31

ABSTRACT

Using comprehensive high-frequency state and local sales tax data, we show that household spending responds strongly to changes in sales tax rates. Even though sales taxes are not observed in posted prices and have a wide range of rates and exemptions, households adjust in many dimensions, stocking up on storable goods before taxes rise and increasing online and cross-border shopping. Interestingly, households adjust spending similarly for both taxable and tax-exempt goods. We embed an inventory problem into a continuous-time consumption-savings model and demonstrate that this seemingly irrational behavior is optimal in the presence of shopping trip fixed costs. The model successfully matches estimated short-run and long-run tax elasticities with a reasonable implied reservation wage of $7-10. We provide additional empirical evidence in favor of this new shopping-complementarity mechanism. While our results reject non-salience of sales tax changes, on average, we also show that upcoming tax changes that are more salient prompt larger responses.

Scott R. Baker
Northwestern University
Kellogg School of Management
Suite 4394
2211 Campus Drive
Evanston, IL 60208
s-baker@kellogg.northwestern.edu

Stephanie Johnson
Northwestern University
Department of Economics
2211 Campus Drive
Evanston, IL 60208
stephgwallace@gmail.com

Lorenz Kueng
Northwestern University
Kellogg School of Management
Suite 4394
2211 Campus Drive
Evanston, IL 60208
and NBER
l-kueng@kellogg.northwestern.edu

1 Introduction

Most economic analysis rests on the assumption that individuals optimize using all publicly available information. Increasingly however, research investigates deviations from this paradigm, allowing individuals to make systematic mistakes or to pay limited attention to relevant information.[1] A range of studies has found empirical evidence of these deviations from frictionless optimizing behavior, in particular in the context of public finance where taxes and fees are often complex (Chetty et al. (2009), Finkelstein (2009), Cabral and Hoxby (2011)).[2] These deviations have profound implications ranging from tax incidence, tax efficiency and optimal taxation (Chetty et al. (2009) and Farhi and Gabaix (2015)) to the effectiveness of countercyclical macroeconomic policies (Gabaix, 2016).

Our first contribution to this literature is a comprehensive analysis of consumer behavior in response to sales tax changes. We study the dynamic response of consumer demand using product-level scanner data and new high-frequency sales tax data at the ZIP code level. Recent research in behavioral public finance has taken a particular interest in state and local sales taxes because of their complexity. Sales taxes in the U.S. are typically not included in 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. Moreover, some goods are exempt from taxation, with different states applying different exemption rules. Furthermore, a wide range of overlapping tax jurisdictions can impose their own sales taxes due to the strong fiscal federalism in the U.S., leading to a large range of tax rates across geographic locations down to the ZIP code level.

Our research design takes advantage of the substantial fiscal implementation lag—the gap between dates when newspapers start covering upcoming sales tax reforms and when sales taxes finally change, which typically occurs months after legislation was passed. We document this fiscal lag with newspaper article counts surrounding sales tax changes. Using Google search results around state sales tax changes, we then show that there is a large and significant spike in searches about sales taxes preceding the sales tax changes, showing that at least some users are using this avenue to find out about upcoming taxes.

We use high-frequency scanner data to show that consumers also act on this information. We find large short-run spending responses to sales tax increases, but small overall spending effects in the long run.[3] These responses are driven by temporary stockpiling of storable and durable goods. Consumers anticipate an upcoming tax increase and bring spending forward to periods with low tax rates, and this form of intertemporal tax arbitrage is more pronounced for more storable and durable goods.

[1]Recent advances in modelling inattention and salience include Sims (2003), Gabaix and Laibson (2006), Chetty, Looney and Kroft (2009), Woodford (2012), Bordalo, Gennaioli and Shleifer (2013), Kőszegi and Szeidl (2013), Gabaix (2014), and Caplin and Dean (2015).

[2]We briefly discuss this literature below.

[3]These responses reflect mostly (compensated) substitution effects, because wealth and income effects occur when consumers learn about the upcoming tax changes, which for forward-looking consumers happens before the taxes change.

These results are characteristic of inventory demand. There are two main benefits from holding inventories of storable consumer goods. First, consumers can shift purchases to periods with low taxes and prices while keeping consumption smooth. Keynes (1936, chp. 13) calls this the *speculative* motive of inventory demand. Second, consumers can reduce shopping fixed costs by holding larger inventories, which Keynes calls the *transaction* motive (e.g., fewer store visits per month).[4] We use transaction time stamps to quantify the importance of this second motive. Consistent with the transaction motive we find that consumers decrease the number of store visits temporarily in the month after a sales tax increase.

Interestingly, we find a very similar spending response for tax-exempt goods. While this behavior seems irrational at first, it is perfectly consistent with optimal inventory management. In fact, this is precisely what one would expect under optimal inventory management with shopping fixed costs! While shopping for taxable goods when tax rates are still low, consumers can reduce the need for future shopping trips by also stocking up on exempt goods.

We derive a novel, parsimoniously parametrized model that relates these results to the recent literature in public finance and macroeconomics mentioned above, which is our second contribution. Business-cycle macroeconomics is interested in the short-run *spending* response, because the effect of countercyclical policies depends on the response of aggregate demand (e.g., Mian and Sufi, 2012).[5] Efficiency cost, static tax incidence, and optimal tax formulas on the other hand depend on long-run *consumption* elasticities (e.g., Farhi and Gabaix, 2015). Our results have implications for both and we use our model to decompose them into the parts that are relevant for each.

For this purpose, we embed an inventory problem with shopping fixed costs into a standard continuous-time consumption-savings model. Households choose an optimal consumption plan and support this plan by managing inventories of two storable consumer goods, a taxable and an exempt good. Optimal inventories trade off holding costs (e.g., depreciation, opportunity costs) against the two main benefits of inventories, the speculative and the transaction motives. The model is successful in matching the observed tax elasticities in the short and long run. We assess the economic magnitude of the responses using the implied reservation wage, which falls in a reasonable range between \$7 and \$10.5.

The model highlights that the observed complementarity of taxable and exempt spending is driven by the shared shopping fixed costs, instead of consumption complementarity. We therefore call this property of demand *shopping complementarity*, because it holds even when the goods are consumption substitutes (which they necessarily are in our two-good model).

We present three additional pieces of evidences to support this new shopping complementarity mechanism. First, we find that only consumers with high shopping fixed costs stock up on both

[4]Keynes identifies a third motive—precautionary demand—in settings with unanticipated price or consumption shocks. See Hendel and Nevo (2006a) for a modern treatment of this idea using scanner data.

[5]Previous research has mostly focused on durable demand, while most of the goods in our data are storable and NIPA would classify them as non-durable.

taxable and exempt goods. Consumers with low fixed costs on the other hand only stock up on taxable, but not on exempt goods. Second, shopping complementarities only apply to goods that are purchased on the same shopping trip. In line with this prediction, we find that online and mail order purchases only increase for taxable goods, because there are no or only few fixed costs to be shared for such purchases (e.g., common search costs).[6] Third, the extent to which exempt spending responds also depends on the degree to which households can bundle their exempt and non-exempt spending. Households that shop at stores that sell solely exempt or solely taxable goods tend to have lower exempt spending responses. Households that typically shop at stores selling a mix of taxable and exempt goods tend to have higher exempt spending responses.

The model yields several additional insights. For instance, it shows that tax and price elasticities of spending can be very different, even with fully rational consumers and full tax salience.[7] The reason is that most changes in (tax-exclusive) posted prices at the store level are unanticipated, while we find that sales tax changes are anticipated well in advance. If goods are storable, then unanticipated price changes prompt starkly different spending responses than anticipated changes. Moreover, sales tax changes are also much more persistent than posted price changes (both temporary sales promotions and "regular" or "reference" price changes). The perceived persistence of a price change also affects the spending response. Intuitively, consumers optimally stock up more in response to an unanticipated temporary price reduction (e.g., store sales) than to an unanticipated permanent price change, since the latter cannot be absorbed with inventory management. Hence, accounting for consumer expectations of tax and price changes is crucial when interpreting observed spending responses.

Another insight from the model is that sales tax changes affect store traffic, because they apply to many goods at once. This extensive margin effect is typically ignored in standard demand estimation, but is important in the context of sales tax changes. Sales tax changes therefore provide a very different source of variation compared to the variation typically used in industrial organization to estimate price elasticities of demand (e.g., product-specific cost or markup shocks). Moreover, we cannot use exempt goods as a control group in a within-store or within-household difference-in-differences setting, since exempt goods are also affected by the anticipated tax change due to the shopping complementarity (a violation of the "stable unit treatment value assumption" in the terminology of program evaluations). In our case, such an approach would bias the estimates all the way to zero.

A concern with our model-based interpretation of the small long-run responses is that it might instead reflect learning-and-forgetting dynamics, i.e., consumers forgetting about sales tax rates over time (Agarwal, Driscoll, Gabaix and Laibson, 2013). This seems reasonable since intertemporal substitution is the main tax avoidance strategy available to most households, and the return of this strategy is limited due to inventory costs. However, we find that consumers

[6]Tax payers are in principle 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.

[7]Coglianese, Davis, Kilian and Stock (2017) make a similar point in the context of excise taxes on gasoline.

change their shopping behavior also in the long run in situations where the return from sales tax arbitrage persists. For instance, consumers who can shop in another tax jurisdiction with lower rates increasingly do so after sales taxes increase in their home ZIP code. This cross-border substitution effect occurs immediately after taxes change and becomes stronger over the next 12 months. Similarly, we find that the effect of a sales tax increase on online spending is also persistent.

Does this mean that sales taxes are fully salient? Of course not. Some consumers are probably not aware of the sales tax rates or changes in such rates, and we suspect that many consumers probably do not know the exemption status of every product they buy, even though they have a general understanding of sales tax exemptions. Our main estimates represent the average response across all households, and there might be substantial heterogeneity in these responses across consumers and by the type of tax change. While we leave a detailed analysis of such heterogeneities for future work, preliminary results for heterogeneity by household characteristics did not yield significant differences in behavior. However, we do find that tax changes that receive relatively more news coverage elicit larger responses (controlling for the size of the tax change). Similarly, spending responses are larger for changes that are driven by (highly-advertised) ballot propositions compared to legislated tax changes that are directly implemented by the state legislature.

Related Literature This study touches on three strands of literature in public economics, macroeconomics, and industrial organization. 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 Chetty et al. (2009), who find that households do not incorporate sales taxes into purchasing decisions and do not respond to changes in sales tax rates, and Finkelstein (2009), who finds that how EZ tolls are paid, via transponder or via cash, affects driving behavior. In a related study, Cabral and Hoxby (2011) find a differential household response to property taxes depending on whether the tax is paid separately or bundled with mortgage payments. There are also important studies that analyze these issues in laboratory experiments; see e.g., Chen, Kaiser and Rickard (2014) and Feldman and Ruffle (2015).

A recent literature in industrial organization focuses on demand dynamics. In a series of papers, Hendel and Nevo (2006a,b, 2013) show that accounting for consumer inventories is crucial when analyzing demand responses to temporary price reductions (i.e., "sales"). Related to these studies, there is a long literature in macroeconomics and finance that estimates the elasticity of intertemporal consumption substitution (EIS), although this literature has not yet reached a consensus.

Our paper contributes to this important literature by highlighting stark differences between the elasticities of intertemporal substitution of spending (EIS-S) and the elasticity of intertemporal elasticity of consumption (EIS-C) when goods are storable. The model identifies the key parameters that drive both responses, showing that a high EIS-S is consistent with a low EIS-C, an insight that goes back at least to Ogaki and Reinhart (1998).

Surprisingly few other studies look at actual sales tax changes in the U.S. A notable exception is Agarwal, Marwell and McGranahan (2017), who study the spending response to sales tax holidays, which are typically well advertised and thus very salient.[8] They find quantitatively similar spending responses. Similarly, recent evidence from other countries is also consistent with our findings. Cashin (2014, 2015) and D'Acunto, Hoang and Weber (2016) examine pre-announced VAT changes in New Zealand, Japan, and Germany, respectively. Value-added taxes (VAT) are included in posted prices and have fewer exemptions than sales taxes in the U.S.

The rest of the paper is organized as follows. Section 2 describes the various data utilized in the paper. Section 3 presents the research design, including the response of newspaper article counts and Google searches for the term "sales tax" around a state tax change. Section 4 estimates the response of household spending and shopping frequency to a pre-announced sales tax increase. Section 5 derives the model. Section 6 provides additional evidence of shopping complementarities. Section 7 discusses additional results, including long-run effects on cross-border and online shopping and the heterogeneity in spending responses as a function of the salience of the tax change. Section 8 concludes.

2 Data

Our empirical analysis uses three types of data: detailed high-frequency household retail spending scanner data at the product-by-store level, monthly combined sales tax rates at the 5-digit ZIP code level, and data that complement our spending-based analysis by providing direct evidence of the fiscal implementation lag, fiscal foresight, and tax salience and its effect on the spending response to sales tax changes.

2.1 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 (e.g., 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 40,000 ZIP codes from 48 states and Washington

[8] Chetty et al. (2009) compare the response of beer sales to excise and sales tax changes at the state level.

DC, excluding Alaska and Hawaii which are not covered by Nielsen's scanner data.

Figure 1 shows the variability in both levels and changes of state and local tax changes. Overall, sales tax changes in our sample period are highly asymmetrical. 80% of total observed changes in sales tax rates are positive, with average total sales taxes increasing from about 6.9% in 2008 to 7.1% in 2014. Restricting to changes in state sales taxes, we find that about 75% of changes are positive, with state sales taxes increasing from 5.4% to 5.7% 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. The average state-level sales tax change (in absolute value) is 0.61% (median 0.5%) and the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles of state level changes are 0.25% and 1%. Local sales tax rate changes are similar both on average (mean 0.54%, median 0.5%) and in their dispersion (standard deviation of 0.37% vs. 0.38% for state level tax changes). 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.

2.2 Retail Spending Data

Household-level retail spending data is obtained from the Nielsen Consumer Panel (NCP, formerly the Homescan Consumer Panel) and store-level retail sales are obtained from the Nielsen Retail Scanner Panel (NRP). The NCP consists of a long-run panel of American households in 52 metropolitan areas from 2004 to 2014. The NPC 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.[9] Broda and Weinstein (2010) provide a more detailed description of the NCP. Einav, Leibtag and Nevo (2010) perform a thorough analysis of the NCP, 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 NCP 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 within a year, so

[9] Approximately 80% of households are retained from year to year.

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, the NCP primarily covers trips to grocery, pharmacy, and mass merchandise stores. 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. To categorize individual products (Nielsen's 'Product Groups') into taxable or tax-exempt goods, we first categorize products into one of the following broad 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. We then assign the 119 Product Groups to these 9 broader exemption categories. For instance, "Crackers", "Dough Products", "Fresh Meat", and "Fresh Produce" would be Product Groups categorized as 'grocery' purchases. Groceries, in turn, are almost always exempt from any state or local sales tax. "Prepared Food Ready to Serve" is assigned to the 'prepared food' category, while "Soft Goods" are treated as 'clothing'. The tax treatment of clothing or prepared food differs by state. Finally, 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 taxable 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.

One concern with the NCP is sample selection since consumers who opt-in to the panel might not be representative based on unobservable characteristics, in particular how much attention they pay to sales taxes. To assess this issue, we also use store-level sales data from the Nielsen Retail Scanner Panel (NRP), which contains price and quantity information of each UPC carried by a covered retailer and spans the years 2006-2014. Nielsen provides the location of the stores at the three-digit ZIP code level (eg. 602 instead of 60208), and we use a population-weighted average sales tax rate using the cross-walk provided by Thomson Reuters.

2.3 Other Data Sources

To obtain direct measures of the fiscal implementation lag, of fiscal foresight by individuals, and 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.[10] We query a set of over 3,000 national, state, and local US newspapers at a monthly frequency from 2008 to 2016.

[10] http://www.newsbank.com/libraries/schools/solutions/us-international/access-world-news.

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 data. This method yields a good match, with only 77 out of 1,468 cities with newspapers being unable to be matched to ZIP codes in our sample.

Google Searches 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.[11] 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 Stephens-Davidowitz (2013), 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

[11] http://www.google.com/trends.

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.

3 Research Design

Our primary empirical methodology is to utilize a difference-in-differences specification with relatively high-frequency spending data at the household/store-month level. The economic interpretation of our empirical results depends crucially on whether the tax changes were anticipated. If tax changes are anticipated, then spending changes and shopping behavior around the tax changes documented in later sections reflect substitution effects, while income and wealth effects take place at the time households learn about an upcoming tax change.

This section therefore provides direct evidence of the long lag between the date when a tax change is legislated and the date when the tax change is finally implemented (i.e., the implementation lag) and of fiscal foresight on the part of individuals. 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. More importantly, using Google Search data we show that some households *actively* acquire information about sales tax rates in advance of the tax rate changes. Both results are consistent with our findings in Section 4 below that households adjust their spending patterns around sales tax changes.

Difference-in-Differences Approach For most of our examination of the impact of changes in sales tax rates, we look at monthly changes in spending and shopping behavior at a household or store level. By construction, the control groups are those households/stores who did not experience a change in the sales tax rate that they face in that month. All regressions include both period and household/store-level fixed effects, thus controlling for seasonal effects, macroeconomic effects, and allowing for household/store-level trends over time.

We estimate the response of various outcomes to changes in total and state sales taxes, respectively, using a difference-in-differences approach by running regressions of the following form:

$$\Delta y_{ht} = \beta \Delta \ln(1 + \tau_{jt}) + \gamma_h + \theta_t + \lambda' z_{ht} + \varepsilon_{ht}. \tag{1}$$

Δy_{ht} is the change in the outcome of interest in month t by household or store h. We consider several outcome variables, including the log of pre-tax expenditures on taxable goods (i.e., expenditures evaluated at posted pre-tax prices), log expenditures on tax-exempt goods, log online and mail-order purchases, the fraction of spending done in a neighboring tax jurisdiction, and measures of shopping frequency. $\Delta \ln(1+\tau_{jt})$ is the log-change of the gross of sales tax rate (since sales taxes are ad valorem) in that month in the corresponding tax jurisdiction j, a zip code or a state (our results are little changed when utilizing the percentage change in sales taxes rather than the logged gross change in prices). γ_h are household/store fixed effects and θ_t are period fixed effects (year and month indicators). z_{ht} are additional time-varying co-variates at the level of the household, store or tax jurisdiction.

Fiscal Implementation Lag Figure 2 shows two metrics of the fiscal implementation lag and of fiscal foresight 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 rate 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. 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. Since this figure is scaled by the size of the change, 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.

Fiscal Foresight While the top panel shows that 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 2 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 the 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 than does the news-based measure. This may reflect a subset of households only realizing sales taxes may have changed over a longer period.

4 Response of Spending and Shopping Behavior

4.1 Taxable Spending

Table 1 shows how retail sales of goods subject to sales taxes change following a change in the sales tax rate. 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 households in the month before the tax increase compared to letting inventories of storable and durable goods deplete in anticipation of a sales tax decrease.[12]

Panel A documents the main results of the analysis. Column 1 shows that following an increase in the combined total sales tax rate of one percentage point (e.g., from 3% to 4%), taxable household retail spending decreases by 2%. This change in spending is measured in the month that the tax change occurs in relative to the moth prior to the tax change.[13] 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).

Panel B tests the robustness of this result. Column 4 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 5 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 6. 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.

Consumers can respond to an anticipated sales tax increase using four main margins of adjustment: moving purchases of storable and durable goods forward to the months before a sales tax increase (and potentially also consumption), shifting spending online and not paying use taxes, shopping in a neighboring tax jurisdiction with a lower sales tax rate, and substitution consumption from taxable to exempt goods. Columns 8 and 9 sequentially shut down the second and third margin (which are the focus of Section 7), thereby restricting the response to intertemporal substitution, while Section 4.2 below analyzes the extent to which households substitute consumption from taxable to exempt goods in the long run. Column 8 restricts the sample to households that did not do any online and mail order purchases in that year, and Column 9 further restricts the sample to households that also do not purchase products in an alternative three-digit ZIP code outside of their own home ZIP code. As one might expect, the point estimates are larger in absolute value for such households that can only respond by engaging in intertemporal substitution, but we cannot reject that the response is the same as in the full sample. The small difference in the response reflects the fact that cross-border and online shopping only makes up a small fraction of purchases for most consumers in the NCP.

As mentioned above, one concern with the NCP is that consumers that select into the panel

[12] However, Column 3 shows that we obtain comparable results when analyzing only tax decreases.

[13] The 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 15$^{\text{th}}$ is the second most common day).

might not be representative based on unobservable characteristics, in particular how much attention they pay to sales taxes. In Panel C we address this issue directly by using store-level data from the Nielsen Retail Scanner Panel (NRP) instead of the household-level NCP data. We find slightly larger responses than in the NCP data, although we again cannot reject the hypothesis that the two estimates are the same because of the larger standard errors of the store-level estimates. However, we would expect larger responses in the store-level data because, while households can take steps to maintain their levels of spending by shifting to untaxed or less-taxed jurisdictions, stores have fewer margins to adjust their exposure to a sales tax increase in the short run.

4.2 Intertemporal Substitution

Any consumer who is aware of an upcoming sales tax increase can at least temporarily avoid paying a higher tax by moving spending forward, even if he does not have the opportunity to shop online or to shop in a neighboring tax jurisdiction with a lower sales tax rate. Hence, intertemporal substitution is the most general adjustment mechanism and this section explores the extent to which consumers take advantage of it.

Intertemporal substitution is the only margin of adjustment that requires consumers to be forward-looking and hence provides a direct test of the salience of upcoming sales tax change changes. Tax incentives for the other two margins of adjustment, cross-border shopping and substitution to online purchases, only change when tax rates change. Estimates of these margins therefore do not test forward-looking behavior.

Shopping Behavior and Intertemporal Substitution Consumers can move both consumption or spending forward to periods with low tax rates. If goods are storable or durable, then these two forms of intertemporal substitution are not the same since consumers can purchases storable goods in advance of the sales tax increase even if they do not change their consumption behavior.[14] If intertemporal substitution of consumption is low, then increases in inventory purchases should decrease the shopping frequency as a larger inventory can support the same consumption rate over a longer period. Hence, comparing the tax elasticity of shopping trips with the tax elasticity of spending provides useful information about the elasticity of intertemporal consumption substitution, a point which the model in Section 5 establishes more formally.

Panel D of Table 1 shows that the number of distinct store visits responds negatively to increases in sales tax rates. The number of trips falls by a similar amount as overall spending, suggesting that the trips adjustment margin is a dominant one. This finding is robust to controlling for household characteristics and for business cycle conditions. Moreover, this extensive margin elasticity is estimated substantially more precisely than the spending elasticity, with standard errors that are 30% smaller.

[14]Again, this analysis focuses on substitution effects that happened after any income and wealth effects have taken place, which typically happens when consumers learn about the upcoming tax increase.

Fiscal Foresight and 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 3 plots the patterns of spending in the months surrounding a change in state sales taxes that is scaled to reflect a 1% tax increase. We find elevated levels of spending in the period preceding a sales tax increase that quickly disappear once the change takes effect. For a sales tax increase of 1%, we see a dramatic fall in spending in the month of the change (period 0) relative to the previous months (period -1), equivalent to a decline in spending of about 2.5%, similar to the more static responses reported in Table 1.

As seen in Figure 3, and across a range of specifications, we find that the average deviation from the household's shopping trend converges to zero relative to the pre-tax-change period. That is, the short-term response of household spending is significantly different than, and greater than, any long-term response. Interestingly, exempt goods see a similar build-up prior to the tax increase and undergo a similar fall in the months afterwards. We study the response of tax-exempt goods to a change in sales tax rates in more detail in Section 4.3 below.

Storability and Intertemporal Substitution Given the 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 or more storable. That is, it would not be feasible to purchase a several-month supply of baked goods 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 manually categorize the 119 product groups contained in the Nielsen data following previous literature on durable goods; see e.g., Cashin (2015).

Second, 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, carbonated beverages, purchased every other week by an average household, would have a value of approximately $\frac{1}{2}$ (2 average purchases per month), while women's fragrances may have a value greater than 12 (purchased less than once a year). 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,

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

$\mathbb{1}\{\text{Trip}_{sg}\}$ 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.

The last two columns of Table 2 show the product groups sorted from highest to lowest number of times a product is purchased in this group per month (or from lowest to highest shopping cycle). The measures of storability and durability correspond fairly well, with durable goods on average having significantly lower purchases per month (0.3) than goods that are not classified as durable (0.8). However, the correlation coefficient is only 0.5 reflecting the fact that many non-durable products are fairly storable. Column 8 shows that the average taxable product is purchased slightly less than once a month suggesting that monthly aggregation of purchases is a sensible choice when analyzing household-level data as in Table 1 rather than product-level data as in Table 2.

We conduct our analysis of heterogeneous behavior across categories of goods at a state-month level by regressing log-changes in pre-tax expenditures on taxable goods in (1) on state sales tax changes interacted with the one of the two measures of product storability and durability, either the continuous average inverse purchase frequency or the discrete hand-classified durability 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.

Table 2 reports results from this analysis. Column 1 shows the baseline spending response to a change in state sales taxes. On average, across categories, we see a decline in spending of approximately 2.5% in the month of a sales tax increase of 1%. This state-level estimate corresponds well with our individual household-level responses of approximately 2.2% (Column 2 of Table 1). Column 2 includes the lead of increases in state level sales taxes to analyze intertemporal substitution patterns. 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. We see that consumers indeed bring spending forward to the month before the sales tax rate increases.

Since there is a lot of heterogeneity in storability and durability of taxable products, this average response is estimated imprecisely. In Panel B we therefore interact the change in sales tax rates with our measures of product durability and storability. Column 3 uses the simple binary indicator for durability while Column 4 uses our continuous measure of storability, the 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 (Column 4). 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.

In Columns 3 and 4, 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.

To show the tremendous heterogeneity in responses by storability, in Column 5 we interact the sales tax changes (both current and lead) with the quartiles of the storability measure. We see an enormous range of demand elasticities. Purchases of products in the top quartile of storability increase by 15% the month before a 1% sales tax increase and then fall by 13% in the month of the tax change. Hence, accounting for a product's storability is very important when price changes are anticipated.

These results are consistent with recent studies of intertemporal household spending behavior. Cashin (2014) 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.

4.3 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.

Table 3 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 1. We find a highly significant response of exempt spending to changes in both state and local sales taxes. This effect remains when limiting the analysis to state sales tax changes and when controlling for household characteristics and 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 3 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.

Overall, the impacts from changes in sales tax changes on tax-exempt goods are similar to those found for taxable goods. The effect is again larger for store-level sales as shown in Panel C (since they are fully exposed to the policy while some consumers can partially avoid the taxes as shown in Section 7), paralleling the larger response of taxable sales in Panel C of Table 1, although even larger in the case of tax-exempt products.

5 A Model of the Shopping Response to Sales Taxes

As we find evidence that complementarities in shopping for both taxable and exempt goods seem to drive household spending, we develop a model of household shopping to determine whether such behavior can be rationalized both qualitatively and quantitatively.[15] In the model, rational forward-looking consumers know about an upcoming sales tax increase at date t_τ and choose consumption continuously. This setting corresponds to our main empirical research design which focuses on the high-frequency spending response around a sales tax increase, but after the upcoming sales tax change has been announced. Hence, these are compensated changes after the direct wealth effect of the higher future tax rate has already been incorporated into the steady-state consumption *levels*. Moreover, because these tax changes occur relatively infrequently within a tax jurisdiction, we use a perfect-foresight model. Consumers face transaction fixed costs of shopping (e.g., time spent shopping, search costs, etc.) and therefore choose discrete-time transaction dates t_n in order to maximize lifetime utility

$$\int_{t=0}^{\infty} e^{-\rho t} u\big(C(t)\big) dt$$

with instantaneous utility $u\big(C(t)\big) = \frac{C^{1-1/\sigma}}{1-1/\sigma}$. $C(t) = \big(\sum_i b_i^{1/\eta} c_i(t)^{1-1/\eta}\big)^{\frac{\eta}{\eta-1}}$ is the composite good consisting of taxable good $c_\tau(t)$ and a tax-exempt good $c_e(t)$, where η is the intra-temporal elasticity of substitution between consumption of taxable and exempt goods. The household receives a constant flow of earned income y and can invest either in a risk-free asset a with continuously compounding return r or in inventories s_i of storable consumer goods which depreciate at rate δ. For notational simplicity and without much loss of generality, assume that the starting date of the analysis, date $t = 0$, is a transaction date. The intertemporal budget constraint is

$$a_0 + \int_{t=0}^{\infty} e^{-rt} y \; dt = \sum_{n=0}^{\infty} e^{-rt_n} K_{t_n}$$

[15] To save space, we provide the more detailed derivation of the model in the Online Appendix.

with initial financial assets a_0 and transversality condition $\lim_{n\to\infty} e^{-rt_n} a_{t_n} = 0$.[16] K_{t_n} are the total shopping trip costs that occur at transaction date t_n and include the fixed costs per trip κ and total consumption expenditures $P_{t_n} S_{t_n}$,

$$K_{t_n}\big(C_{t_n}, \Delta t_n\big) = \kappa + P_{t_n} S_{t_n}\big(C_{t_n}, \Delta t_n\big). \tag{2}$$

C_{t_n} denotes instantaneous consumption of the composite good at the beginning of the endogenously chosen period of length $\Delta t_n = t_{n+1}^- - t_n$, where t_{n+1}^- is the limit from below of the next transaction date t_{n+1}.

S_{t_n} is the necessary beginning-of-period inventory of the composite good to support consumption until the next transaction date. Inventory $s_i(t)$ of good i solves the differential equation $\dot{s}_i(t) = -\delta s_i(t) - c_i(t)$ at any point during the shopping interval. Because there is no uncertainty, there is no precautionary inventory demand and households optimally exhaust inventories fully right before the next shopping trip, $s_i(t_{n+1}^-) = 0$. This terminal condition determines the necessary initial inventory at the beginning of the shopping cycle. The expenditure-minimizing cost of a unit of the composite consumption good purchased at date t_n is $P_{t_n} = \big(\sum_i b_i p_{i,t_n}^{1-\eta}\big)^{1/(1-\eta)}$ and Hicksian demand is $c_{i,t_n} = b_i(\frac{p_{i,t_n}}{P_{t_n}})^{-\eta} C_{t_n}$. Total expenditures at the beginning of the shopping interval can be expressed in terms of inventory of the composite good, $P_{t_n} S_{t_n} = \sum_i p_{i,t_n} s_{i,t_n}$, with $s_{i,t_n} = b_i(\frac{p_{i,t_n}}{P_{t_n}})^{-\eta} S_{t_n}$. Defining

$$f(\Delta t_n; \alpha) = \int_{t=0}^{\Delta t_n} e^{\alpha t} dt = \frac{e^{\alpha \Delta t_n} - 1}{\alpha}$$

the consumer purchases $S_{t_n} = C_{t_n} \cdot f(\Delta t_n; \phi)$ at date t_n given price index P_{t_n}, where $\phi = \delta - \sigma(\delta + \rho)$ is the effective discount rate of utility over the transaction interval Δt_n, which accounts for the pure time preference, the user cost of inventory, and the consumer's willingness to shift consumption intertemporally within a period between two shopping transactions.

We formulate this problem as a dynamic program by discretizing the consumption plan to match the shopping intervals, building on an early model by Howitt (1977). For this purpose, we define the indirect utility function of consumption between shopping transactions as

$$\begin{aligned} U(C_{t_n}, \Delta t_n) &= \max_{\{C(t)\}} \Big\{ \int_{x=0}^{\Delta t_n} e^{-\rho x} u\big(C(t_n + x)\big) dx : \int_{x=0}^{\Delta t_n} e^{\delta x} C(t_n + x) dx = S_{t_n} \Big\} \\ &= u(C_{t_n}) \cdot f(\Delta t_n; \phi). \end{aligned} \tag{3}$$

Defining total wealth $w_{t_n} = a_{t_n} + y/r$, the problem can be written as a dynamic program,

$$V(w_{t_n}) = \max_{C_{t_n}, \Delta t_n} \Big\{ U(C_{t_n}, \Delta t_n) + e^{-\rho \Delta t_n} V(w_{t_{n+1}}) : \ w_{t_{n+1}} = e^{r \Delta t_n}(w_{t_n} - K_{t_n}) \Big\}. \tag{4}$$

[16] We write continuous time variables as $x(t)$ and discrete time variables as x_{t_n}, where $x_{t_n} = x(t_n)$ is the value of x at the beginning of the endogenously chosen transaction interval.

The envelope theorem requires that an additional dollar received at the beginning of each transaction period has the same present utility value,

$$V'_{t_n} e^{-r\Delta t_n} = e^{-\rho \Delta t_n} V'_{t_{n+1}}. \tag{5}$$

Optimal consumption and beginning-of-period inventory of the composite good are characterized by the first-order condition

$$\partial_C U'_{t_n} = \partial_C K'_{t_n} \cdot V'_{t_n}. \tag{6}$$

$\partial_C K'_{t_n} = P_{t_n} f(\Delta t_n; \phi)$ is the effective price of consumption taking into account the inventory costs. Combining equations (5) and (6) we obtain the familiar Euler equation for the growth rates of unobserved beginning-of-period consumption,

$$\frac{C_{t_{n+1}}}{C_{t_n}} = e^{\sigma(r-\rho)\Delta t_n} \left(\frac{P_{t_{n+1}}}{P_{t_n}}\right)^{-\sigma},$$

and of observable beginning-of-period inventory of the composite good,

$$\frac{S_{t_{n+1}}}{S_{t_n}} = \frac{C_{t_{n+1}}}{C_{t_n}} \cdot \frac{f(\Delta t_{n+1}; \phi)}{f(\Delta t_n; \phi)}. \tag{7}$$

The less familiar necessary condition determining the optimal transaction interval is

$$\partial_{\Delta t} U'_{t_n} - \partial_{\Delta t} K'_{t_n} \cdot V'_{t_n} = e^{-\rho \Delta t_n} \Big[\rho V_{t_{n+1}} - r w_{t_{n+1}} \cdot V'_{t_{n+1}} \Big]. \tag{8}$$

The left-hand side captures the net marginal utility at date t_n of increasing the time until the next shopping transaction at date t_{n+1}, which equals the present value of the additional consumption utility net of the additional cost to support the extension of the transaction interval. The terms in square brackets on the right-hand side capture the net marginal cost from starting the next period later, which equals the cost from delaying the continuation value net of the additional interest earned.

Steady State From equation (5) we see that the stationary state, which starts at the first transaction date t_{ss} after the tax increase, requires $r = \rho$ unless the value function is linear. The optimal inventory and transaction intervals in the stationary state are jointly determined by combining equations (4) to (6),

$$(1-\sigma) \frac{\kappa}{P_{t_{ss}} S_{t_{ss}}} = e^{\phi \Delta t_{ss}} \frac{f(\Delta t_{ss}; r)}{f(\Delta t_{ss}; \phi)} - 1, \tag{9}$$

and by the budget constraint in the stationary state $w_{t_{ss}} = (1 - e^{-r\Delta t_{ss}})^{-1} (P_{t_{ss}} S_{t_{ss}} + \kappa)$. In the stationary state, the consumer trades off the additional user cost by marginally extending the shopping trip interval against the marginal benefit of pushing the fixed costs further into the future. To gain intuition, we can related this condition to the familiar square-root formula from static inventory models *if* we assume that the consumer is unwilling to substitute consumption intertemporally ($\sigma = 0$) and if we take a second-order approximation of (9) around $\Delta t_{ss} = 0$,

$$\Delta t_{ss} \approx \sqrt{\frac{\kappa}{\frac{\delta+r}{2} P_{t_{ss}} C_{t_{ss}}}}.$$

Higher transaction fixed cost as a fraction of total spending per trip lead to less frequent shopping, while higher user costs (depreciation and forgone interest) lead to more frequent shopping. However, in general with $\sigma > 0$ this is not a good approximation.

5.1 Shopping Response to an Anticipated Sales Tax Increase

Figure 4 shows the evolution of composite consumption and inventories (left y-axis), the increase in the price index due to an anticipated sales tax increase (right y-axis) and endogenously chosen transaction intervals (x-axis). To make this example as stark as possible, we use a very large tax change of 10% and an unrealistically large elasticity od intertemporal substitution of 6 and a low fixed cost of \$2. Both choices are made only for this figure.

Because the sales tax increase is fully anticipated by forward-looking consumers, the problem of choosing consumption, inventories, and transaction dates is non-stationary and standard inventory models cannot easily be applied to this setting. However, as the figure shows, we can divide the solution into three stages: (i) the pre-period shopping intervals (Δt_{ss-q} for $q \geq 2$) that occur while the consumer faces the old sales tax rate in the current and the next shopping trip, (ii) the interim shopping interval Δt_{ss-1}, which is the last trip before the tax increase, and (iii) the final stationary state of shopping intervals Δt_{ss} that occur under the new sale tax rate.

Tax Elasticities Next we derive analytic expressions for the tax elasticities of the key variables in the model—consumption, inventory, and shopping intervals—and map them to the observed spending and trip elasticities estimated in Section 4. The consumption Euler equation governs the wealth-compensated consumption elasticity to an anticipated sales tax increase at the time of the tax change (i.e., after the consumer has been informed about the new tax rate and hence after the wealth effect on the steady-state consumption level has been incorporated),

$$\begin{aligned} \varepsilon_{c_i} \equiv \frac{d\ln\big(c_i(t_{ss})/c_i(t_{ss-1})\big)}{d\ln(1+\tau_{t_{ss}})} &= -(\sigma-\eta)\frac{d\ln(P_{t_{ss}}/P_{t_{ss-1}})}{d\ln(1+\tau_{t_{ss}})} - \eta\frac{d\ln(p_{i,t_{ss}}/p_{i,t_{ss-1}})}{d\ln(1+\tau_{t_{ss}})} \\ &= -(\sigma-\eta)B_\tau - \eta\ 1_{\{i=\tau\}}. \end{aligned} \tag{10}$$

The second line uses $d\ln(P_{t_{ss}}/P_{t_{ss-1}})/d\ln(1+\tau_{t_{ss}}) = B_\tau$, where $B_\tau = p_{\tau,t_{ss}} s_{\tau,t_{ss}}/(P_{t_{ss}} S_{t_{ss}})$ is the expenditure share of taxable goods in steady state. For taxable goods, the reduced-form consumption elasticity is unambiguously negative (or non-positive), while the sign of the elasticity of tax-exempt consumption depends on the relative size of the two structural elasticities, the *inter*temporal substitution elasticity σ and the *intra*temporal substitution elasticity η.

The two reduced-form consumption elasticities ε_{c_τ} and ε_{c_e} are not directly observable in the data since consumption differs from spending. Moreover, substitution of consumption is not the only or even the main adjustment margin available to consumers. Instead, consumers can also respond by bringing spending forward and thereby extending the time until the first transaction under the higher tax rate. Using (7) and (9), the increase in the length of the interim shopping interval Δt_{ss-1} (and hence the start date t_{ss} of the steady state under the higher tax rate) can be expressed in closed form as

$$\Delta t_{ss-1} - \Delta t_{ss} = \frac{\ln(P_{t_{ss}}/P_{t_{ss-1}})}{\delta + r}.$$

The consumer trades off the marginal return from bringing spending forward before taxes increase (numerator) against the additional user cost incurred during the shopping interval (denominator). The shopping interval does not change unless the price level is expected to change. Therefore, the structural interpretation of the short-run elasticity of the shopping trip interval is

$$\varepsilon_{\Delta t_{ss-1}} \equiv \frac{d\ln(\Delta t_{ss}/\Delta t_{ss-1})}{d\ln(1+\tau_{t_{ss}})} \approx -\frac{B_\tau}{(\delta + r)\Delta t_{ss}}. \tag{11}$$

Combining (7) with (11), the short-run elasticities of taxable and exempt spending are

$$\varepsilon_{s_{i,t_{ss-1}}} \equiv \frac{d\ln(s_{i,t_{ss}}/s_{i,t_{ss-1}})}{d\ln(1+\tau_{t_{ss}})} \approx \varepsilon_{\Delta t_{ss-1}} + \varepsilon_{c_i} \tag{12}$$

and the long-run spending elasticities are well approximated by the corresponding consumption elasticities,

$$\varepsilon_{s_{i,\infty}} \equiv \frac{d\ln(s_{i,t_{ss}}/s_{i,t_{ss-q}})}{d\ln(1+\tau_{t_{ss}})} \approx \varepsilon_{c_i} \tag{13}$$

for $q \geq 2$.

Calibration and Aggregation We calibrate the model to match steady state values, the long-run responses, and one short-run elasticity: the fall of taxable spending in the month of the tax increase relative to steady state. We then assess the model along two dimensions. First, we analyze whether the model can generate the short-run spending dynamics for both taxable and tax-exempt goods shown in the left panel of Figure 3. Second, while matching the short-run dynamics is an important test for the model to pass, it does not tell us whether the magnitudes of the observed responses are economically reasonable. To answer this question, we calculate the revealed reservation wage implied in the shopping fixed costs necessary to match the short-run spending dynamics.

The long-run response of tax-exempt spending is close to zero as seen in Figure 3. Looking at equations (10) and (13), this implies that intertemporal and intratemporal consumption substitution elasticities σ and η must be of similar size. To assess whether both are small or large, we calculate the relative difference between the long-run spending elasticities of taxable and tax-exempt, which equals the intratemporal consumption elasticity. The estimated difference is small and hence we set both elasticities to 0.29, although this estimate is not statistically different from zero (standard error of 0.48).

The two remaining parameters calibrated to steady state values are the share parameters b_i and the steady-state transaction interval Δt_{ss}. The share parameters are set to match the average expenditure shares of taxable and tax-exempt goods in months without sales tax changes: 0.55 and 0.45, respectively. We set the steady state interval to 0.27 months, corresponding to the sample average of 8.3 days between two shopping trips. The effective annual risk-free rate is 3% and the depreciation rate is set such that the steady state equation (9) holds. Relative pre-tax prices are normalized to one.

Before calibrating the remaining parameter κ to match the decline in taxable spending in the month of the tax increase, we need to adjust for the fact that model time is measured at trips frequency rather than monthly frequency (which did not matter for steady state calibrations). To eliminate artificial lumpiness resulting from the allocation of trips to months,[17] we distribute shopping start dates of a unit mass of otherwise identical consumers uniformly on an interval that starts at date 0 and has length Δt_{ss-2}.[18] We can then aggregate per trip quantities s_{i,t_n} (valued at pre-tax prices normalized to 1) to monthly spending quantities $s_{i,t}$, where t measures event time, as follows:

$$\begin{aligned}
s_{i,-2} &= \frac{s_{i,t_{ss-2}}}{\Delta t_{ss-2}} \\
s_{i,-1} &= (1-\Delta t_{ss-2})\frac{s_{i,t_{ss-2}}}{\Delta t_{ss-2}} + s_{i,t_{ss-1}} \\
s_{i,0} &= [1-(\Delta t_{ss-1}-\Delta t_{ss-2})]\frac{s_{i,t_{ss}}}{\Delta t_{ss}} \\
s_{i,1} &= \frac{s_{i,t_{ss}}}{\Delta t_{ss}}.
\end{aligned}$$

In month -1, the rate of spending over the initial fraction $(1-\Delta t_{ss-2})$ of the month is $s_{i,t_{ss-2}}/\Delta t_{ss-2}$ (spending per trip $s_{i,t_{ss-2}}$ times shopping frequency $1/\Delta t_{ss-2}$). The rate of spending over the remaining fraction Δt_{ss-2} of the month is $s_{i,t_{ss-1}}/\Delta t_{ss-2}$. In month 0, the rate of spending is initially 0 on an interval of length $\Delta t_{ss-1} - \Delta t_{ss-2}$ because each household stocked up more on interim shopping trip t_{ss-1} in month -1. The spending rate on the remaining fraction of month

[17] To see the issue arising from time aggregation of a single consumer, suppose that $\Delta t_{ss} = \Delta t_{ss-2} = 0.25$, which holds approximately in the data, but $\Delta t_{ss-1} = \Delta t_{ss} + \varepsilon$ for an arbitrarily small $\varepsilon > 0$ due to additional stockpiling. This consumer would therefore only make 3 trips in month -1 but 4 trips in any other month, which would lead to very large monthly spending changes even though spending per trip in the continuous-time model would increase only very little on trip t_{ss-1} relative to all other trips (i.e., $s_{i,t_{ss-1}} \approx s_{i,t_n}\ \forall\ t_n \neq t_{ss-1}$).

[18] Equivalently, we could endow consumers with different inventory levels at date 0.

0 is $s_{i,t_{ss}}/\Delta t_{ss}$.[19]

With this aggregation, setting $\kappa = \$5.2$ matches the decline of 1.45% in taxable spending in the month of the tax increase relative to the steady state (left panel of Figure 3). These fixed costs per trip enter the model solution as a fraction of total spending per trip, which we set to its sample average of $P_{t_{ss}} S_{t_{ss}} = \$83$.[20]

Model Evaluation The right panel of Figure 3 shows that the model is successful in producing the short-run spending patterns surrounding a 1% sales tax increase, both qualitatively and quantitatively. The simple model fits the data well even though it only uses three data points for the calibration (the long-run spending changes of 0 and -0.29 and the short-run deviation of taxable spending from the steady state of -1.45).

We assess the economic magnitude of these responses using auxiliary data from the American Time Use Survey (ATUS) to calculate the reservation wage implied in the estimated fixed costs. Across our sample period, people in the ATUS report spending on average 0.1 hours per day on grocery shopping. The average and median number of days between two trips to a grocery store by consumers in the Nielsen data are 6 and 4, respectively, implying about half an hour spent in the grocery store per trip. Hamrick and Hopkins (2012) estimate that the average respondent spends another 15 minutes on travel per grocery shopping trip. These estimates result in an economically reasonable range for the reservation wage between \$7 and \$10.5.[21]

6 Shopping Complementarity

The reduced-form response of tax-exempt spending in Section 4.3 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,

[19] Technically, starting with a unit mass of consumers that have completely unsynchronized shopping cycles, a systematic change in long-run shopping intervals Δt_{ss} relative to Δt_{ss-2} in response to a sales tax increase introduces small echo effects in aggregate monthly steady state spending either due to a recurring hole in the distribution of shoppers over a small interval of length $|\Delta t_{ss-2} - \Delta t_{ss}|$ or a small amount of excess mass (bunching) over an interval with the same length. However, this difference is very small in practice such that our aggregation procedure yields a very good approximation of average monthly spending dynamics.

[20] We exclude small transactions from this calculation in order to identify shopping trips rather than say lunch purchases for immediate consumption for example, which are not captured by the model.

[21] Note that this reservation wage is after taxes, i.e., after income and payroll taxes.

taking advantage of the detailed information on consumer and retailer locations. First, 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. This may reflect both higher reservation wages as well as higher direct costs like gasoline or public transit fees. Second, we test the extent to which the exempt spending response depends on the degree to which consumers can bundle their exempt and non-exempt spending. Intuitively, consumers that shop at stores that sell solely exempt or solely taxable goods should have lower exempt spending responses, while consumers that typically shop at stores selling a mix of taxable and exempt goods should have higher exempt spending responses. 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 4 presents the results of this analysis. Panel A investigates 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 1 and 2 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 2.2%. 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.

Panel B examines the response of exempt spending across two different types of households. Using the granular Nielsen purchase data, we determine how skewed a given household's average shopping trip is towards either exempt or taxable purchasing:

$$\text{Trip Complementarity}_i = 1 - \frac{\sum_j |(T_{ij} - 0.5)| \times 2}{\sum_j \mathbf{1}}.$$

That is, if all of household i's trips (indexed by j) are for 100% taxable ($T_{ij} = 1$) or 100% exempt goods ($T_{ij} = 1$), his average trip complementarity measure would receive a value of 0. If each trip was composed of 50% taxable goods and 50% exempt goods ($T_{ij} = 0.5$), the measure would take on a value of 1.

In Column 5, we look at the highest quartile of households along this measure. We find that exempt spending for this group responds strongly to changes in sales taxes. In contrast, the quartile of households whose taxable and exempt spending is conducted at largely different stores sees a much smaller (and insignificant) change in spending on exempt goods shown in Column 6.

Finally, Panel C tests for the asymmetric response of online spending and mail order purchases 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 (Column 7), while spending on taxable goods increases significantly (Column 8). This response suggests that consumers evade sales taxes by substituting to online platforms and failing to pay use taxes on those purchases.

Identification of Demand in the Presence of Shopping Complementarities 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.

7 Additional Evidence of Optimizing Behavior

In this section, we first inspect additional mechanisms that might prompt responses to a sales tax increase, in particular some tax incentives that persist after a sales tax change. We then study 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. Consistent with previous research, we find that more salient tax changes prompt larger spending responses.

7.1 Persistent Tax Incentives Prompt Long-Run Responses

Section 4 has shown that intertemporal substitution is only a temporary response and spending quickly reverts back to pre-change levels. In the model, this happens because the intertemporal consumption substitution elasticity and the intratemporal substitution elasticity between taxable to tax-exempt products are both low. An alternative explanation of the small long-run responses is that it might instead reflect learning-and-forgetting dynamics, i.e., consumers forgetting about sales tax rates over time. We test this explanation by analyzing the spending response to changes in tax incentives that can be exploited more persistently. Specifically, we first test whether consumers who can shop in another tax jurisdiction with lower rates increasingly do so after a sales tax increase in their home ZIP code, and whether this response persists in the long run. We then extend our analysis of online shopping in Section 6 to test whether the effect of a sales tax increase on online spending is also persistent. Table 5 presents the results from this analysis.

Jurisdictional Tax Avoidance: '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 'alternative' three-digit ZIP code or state (outside one's 'home' ZIP code or state).

Panel A analyzes such cross-border shopping behavior. Column 1 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 estimate (1) by regressing the fraction of a household's total spending in an alternative three-digit ZIP code outside its own home ZIP code (i.e., residential five-digit ZIP code) on log-changes in the gross sales tax rate interacted with the average alternative three-digit ZIP spending over all household-months. Column 2 shows precisely this mechanism, a significant increase in alternative-ZIP spending in the month of the tax change for households who had already been conducting some of their shopping in alternative three-digit ZIP codes. Columns 3 and 4 show the same pattern for purchases across state lines, which are less common (average spending share of less than 2% compared to the 8% spending share in alternative three-digit ZIP codes). This signals that, for households who could conceivably substitute spending into a different tax jurisdiction, an increase in the sales tax in their residential tax jurisdiction makes them shift additional spending to that alternative tax jurisdiction.

The short-run month-to-month behavior is estimated only imprecisely and the interaction terms are not statistically significant. One reason is that in contrast to the significant intertemporal substitution of short-run spending shown in Section 4, consumers should rationally only increase cross-border spending *after* sales taxes have increased in their home ZIP code, not before. Since most products in the Nielsen data are fairly storable or durable, the next cross-border shopping trip might be months away such that this adjustment occurs only gradually over time. Columns 5 and 6 show that the long-run response is indeed larger and more precisely estimated. The interaction effects are now two to three times larger than in the short-run and statistically significant.

An important note about substitution across jurisdictions is that while 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 Davis, Knoepfle, Sun and Yannelis (2016) also looks at the geographical substitution patterns surrounding sales taxes. Using 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, they estimate 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, another persistent adjustment margin to which turn next.

Use Tax Evasion: 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 if the merchant does not maintain a physical presence in the same state as the purchaser. During our sample period, most 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, as mentioned before 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 sales taxes increase.

Fortunately, the Nielsen Consumer Panel data categorizes purchases made from online merchants separately from brick-and-mortar retailers. Panel B shows the result of this analysis. We find that household shift spending to these online merchants in the month following a sales tax increase (Column 7) and this substitution persists in the long run (Column 8).[22]

These coefficients suggest that online spending by 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, Einav, Knoepfle, Levin and Sundaresan (2014) find an online-offline substitution elasticity of 1.8, which is in line with our estimate of 1.6. Similarly, Baugh, Ben-David and Park (2015) 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.

7.2 Tax Salience and Announcement Effects

A natural question that arises given the results displayed in Figure 2 is whether tax salience plays an additional role in consumers response to a sales tax rate chang and whether news about

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

future sales tax changes prompt a response via an income or wealth effect. While the results in Section 4 document a significant degree of tax 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 is not fully salient (e.g., cookies vs. candies). In this section, we 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). Table 6 presents the results from this analysis.

Tax Salience: Evidence from Ballot Initiative Panel A uses two measures of tax salience and examines 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 interact this news-based measure with changes in state sales taxes. To facilitate the quantitative interpretation, we normalize the news measure by its standard deviation. Since it is a residual, the resulting transformation has mean zero and a unit standard deviation (i.e., a standard score). We again find that, in general, sales tax changes have a negative relationship with spending in the month of the tax change, comparable with the baseline effects reported in Column 2 of Tables 1 and 3. Moreover, changes that had more news coverage (conditional on the size of the change) also had larger declines. The coefficient on the interaction term of Column 1 shows that an increase in news coverage of one standard deviation would increase the spending response to a 1% sales tax by about 20% (from -1.8% to -2.1%). The effect is again similarly shared by taxable and tax-exempt spending.

Columns 4 to 6 take a different approach to testing heterogeneity in household responses across sales tax changes with different salience. Here we utilize data on state-level ballot measures that changed state sales taxes. Our prior is that sales tax changes 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. Consistent with this hypothesis, 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.

'News' Response: Income, Wealth and Substitution Effects Panel B of Table 6 demonstrates 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 state 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.[23]

In a model with fully informed and rational consumers, households would perceive this future tax increase as a persistent increase in future prices. At the time of the announcement (which is before time 0 in the model of 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 provides suggestive evidence that this effect might play a role, on average, across all ballots (whether they passed or failed), with the act of voting on the ballot being associated with a 0.5% decline in household retail spending. We further refine the analysis by separating these ballots into those that failed and those that passed, finding opposite signed coefficients. 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. These results are consistent with forward-looking behavior on the part of consumer, although they are not statistically significant.

8 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 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 do respond to changes in sales taxes, both following a 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 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 on average aware of other methods to avoid sales taxes,

[23] Ideally 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.

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 consumers respond to sales tax changes along many dimensions suggested by economic theory, we also demonstrate that increases in the amount of information presented to consumers about upcoming changes to sales taxes tends to induce larger spending responses. These results imply that the salience of taxes is an important determinant of how households respond to taxes.

References

Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson, "Learning in the Credit Card Market," *Working Paper*, 2013.

_ , Nathan Marwell, and Leslie McGranahan, "Consumption Responses to Temporary Tax Incentives: Evidence from State Sales Tax Holidays," *American Economic Journal: Economic Policy*, 2017.

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," *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.

Coglianese, John, Lucas W. Davis, Lutz Kilian, and James H. Stock, "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand," *Journal of Applied Econometrics*, 2017, *32* (1), 1–15.

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?," *QME*, 2010, *8* (2), 207–239.

Farhi, Emmanuel and Xavier Gabaix, "Optimal Taxation with Behavioral Agents," *Working Paper*, 2015.

Feldman, Naomi E. and Bradley J. Ruffle, "The Impact of Including, Adding, and Subtracting a Tax on Demand," *American Economic Journal: Economic Policy*, 2015, *7* (1), 95–118.

Finkelstein, Amy, "EZTax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 2009.

Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

_ , "A Behavioral New Keynesian Model," *Working Paper*, 2016.

_ and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Hamrick, Karen S. and David Hopkins, "The time cost of access to food–Distance to the grocery store as measured in minutes," *International Journal of Time Use Research*, 2012, *9* (1), 28–58.

Hendel, Igal and Aviv Nevo, "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 2006, *74* (6), 1637–1673.

_ and _ , "Sales and Consumer Inventory," *RAND Journal of Economics*, 2006, *37* (3), 543–561.

_ and _ , "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, 2013, *103* (7), 2722–2751.

Howitt, Peter, "Intertemporal Utility Maximization and the Timing of Transactions," *American Economic Review*, 1977, *67* (2), 156–165.

Keynes, John Maynard, *The General Theory of Employment, Interest and Money*, Macmillan Cambridge University Press, 1936.

Kőszegi, Botond and Adam Szeidl, "A model of focusing in economic choice," *The Quarterly Journal of Economics*, 2013, *128* (1), 53–104.

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.

Ogaki, Masao and Carmen M. Reinhart, "Measuring Intertemporal Substitution: The Role of Durable Goods," *Journal of Political Economy*, 1998, *106* (5), 1078–1098.

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.

Woodford, Michael, "Inattentive Valuation and Reference-Dependent Choice," *Working Paper*, 2012.

Figure 1 – ZIP Code Level Sales Tax Rates

(a) Sales tax rates

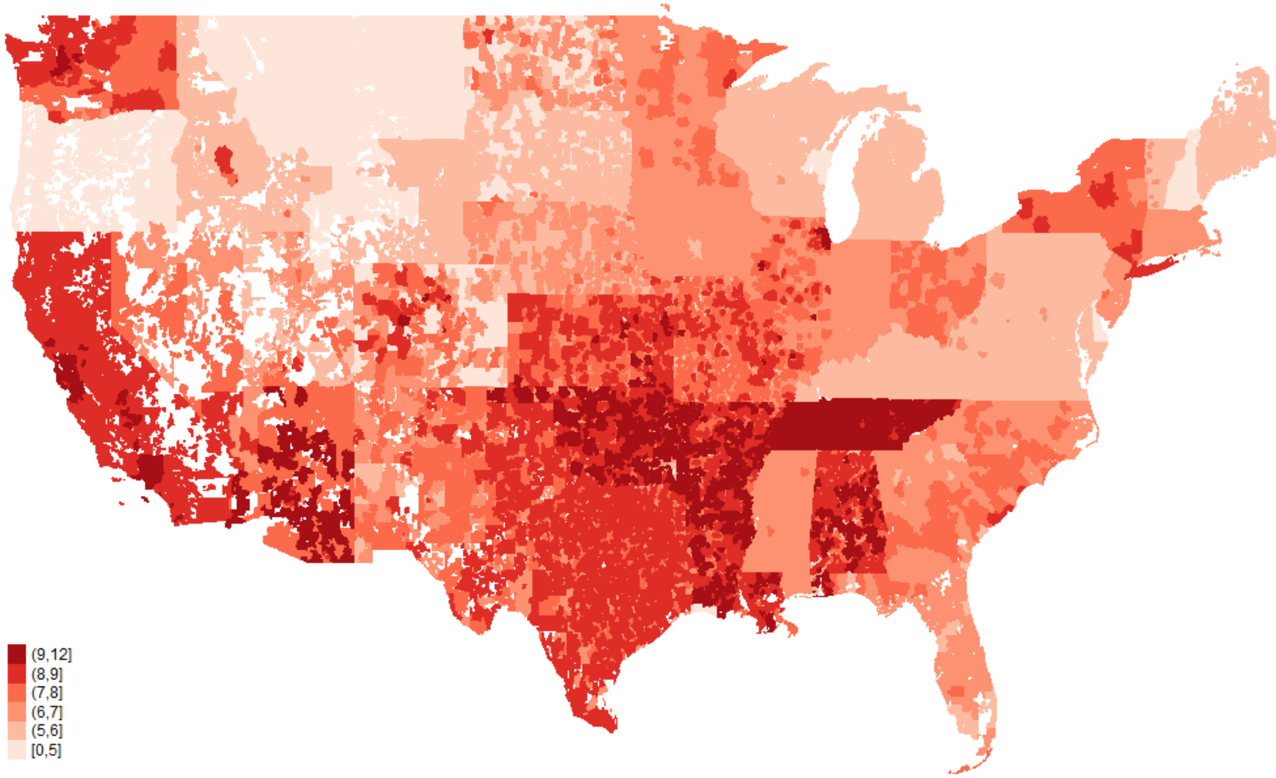

(b) Tax rate changes, 2008-14

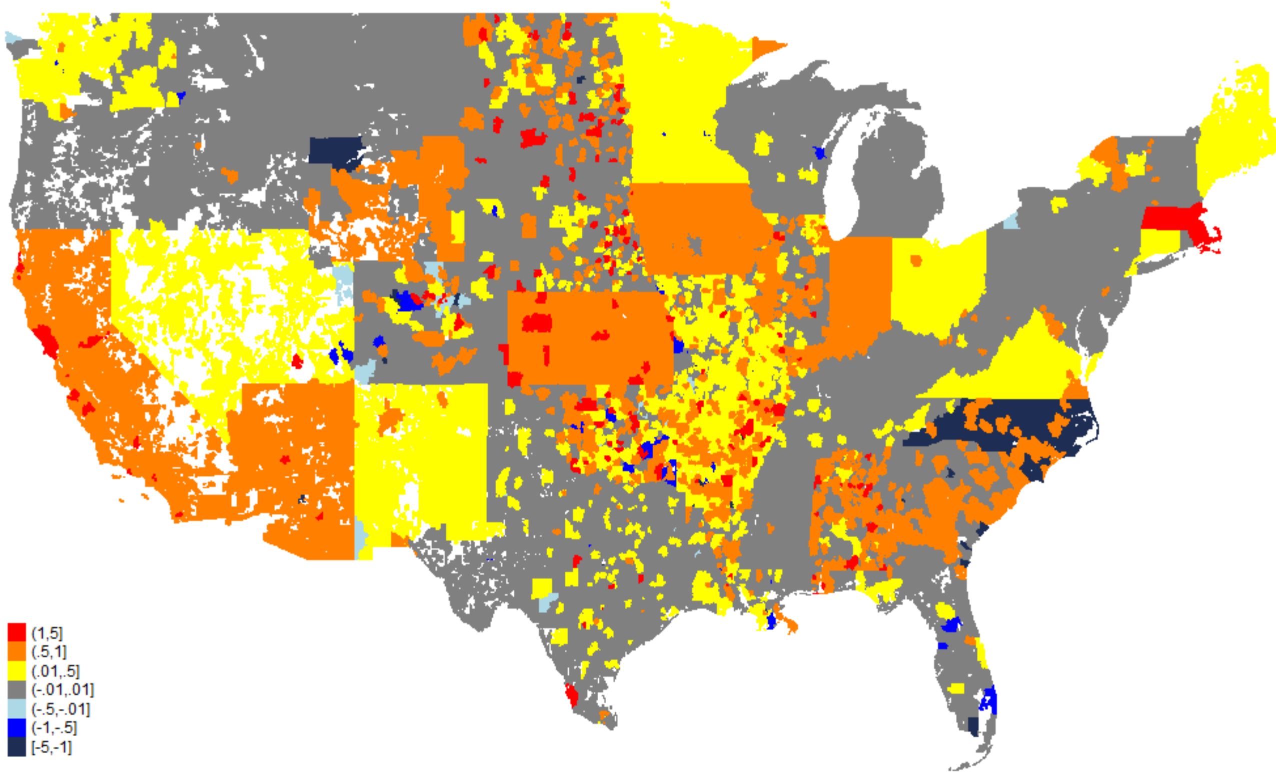

Notes: Maps plot the maximum level (a) respectively change (b) of total sales tax rates in each five-digit ZIP code for years 2008-14, matching the sample period of the Nielsen Consumer Panel. Sales tax rates are expressed in percentages. Total sales tax rate changes may be driven by changes in state,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

Figure 2 – Response of Newspaper Coverage and Google Searches to State Tax Changes

(a) News article ratio around state sales tax rate changes (in %)

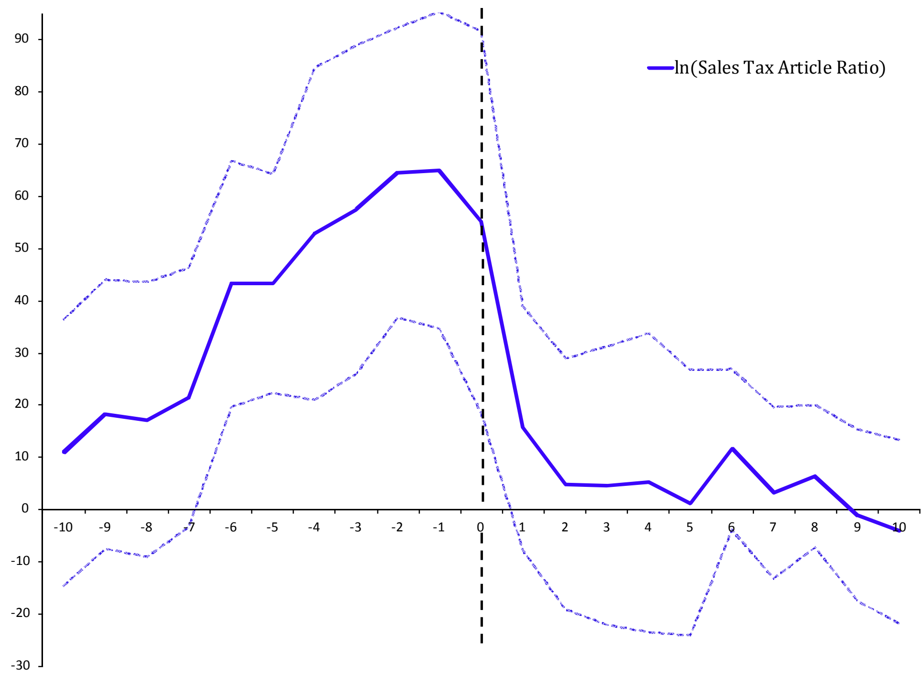

(b) Google Search around state sales tax rate changes (in %)

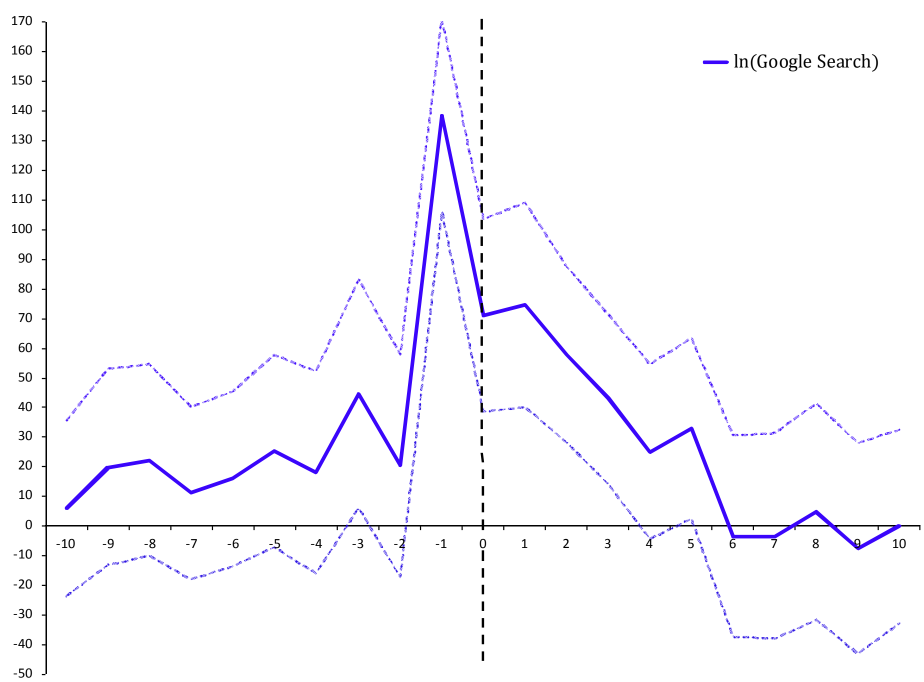

Notes: Top panel plots coefficients from a regression of the ratio of news articles that contain the term 'sales tax' or 'sales taxes' as a fraction of all newspaper articles in a given month across newspapers in that state. Y-axis units are percentage points. News articles taken from Access World News 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 3 – Spending Dynamics around a Sales Tax Increase: Estimation and Model

(a) Estimation: β coefficients from log-level regression

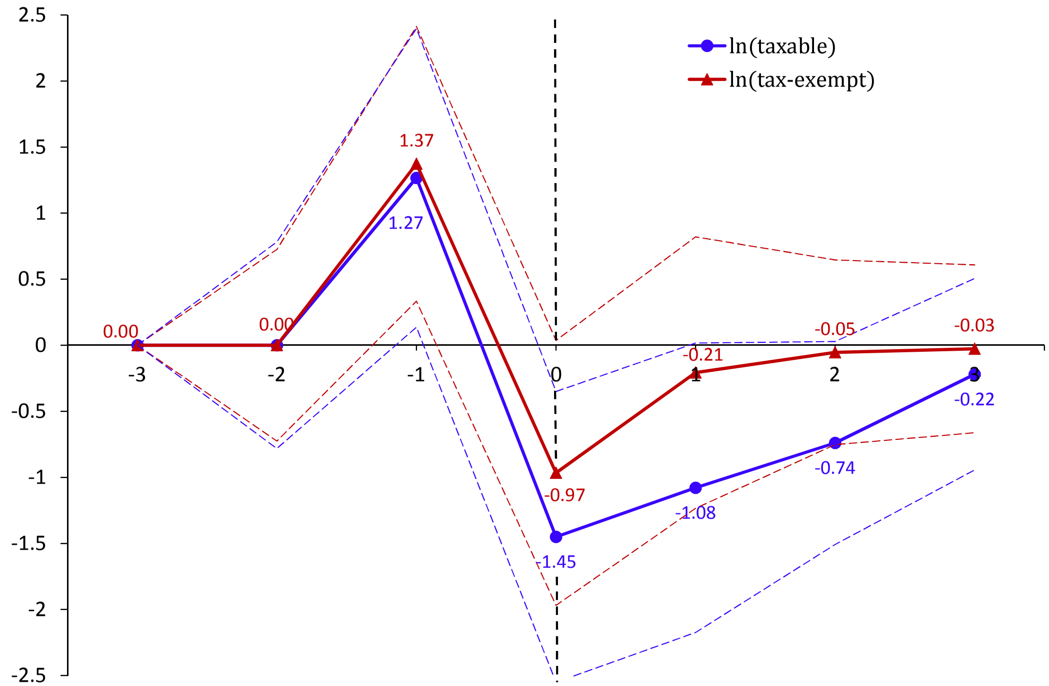

(b) Model: Log-levels in discrete time (monthly)

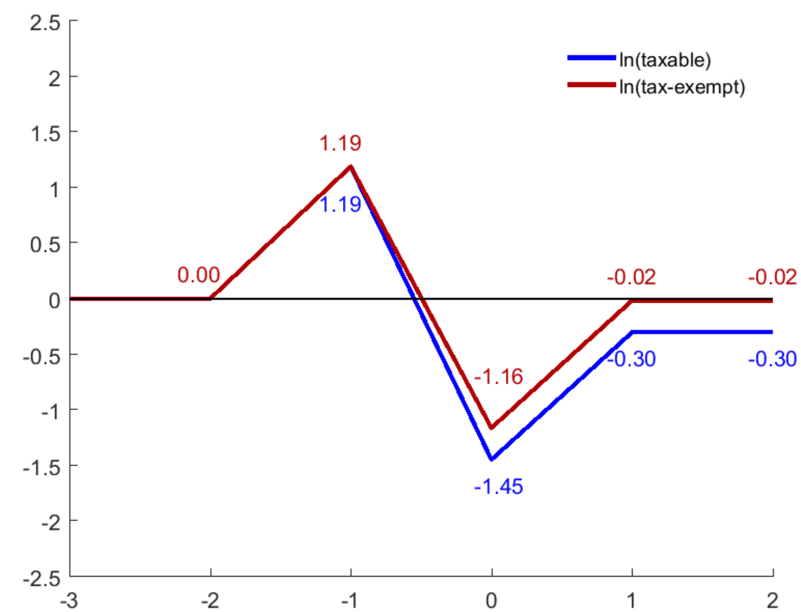

Notes: Left panel plots coefficients of a regression of the logged amount of pre-tax household retail spending on taxable and exempt products on leads and lags of total sales tax rate increases. All coefficients are scaled to an increase in sales taxes of 1%. Dashed lines represent 95% confidence intervals from standard errors clustered at the zipcode level. Periods -1, 0 and 1 reflect the three months around the tax increase and periods -2, 2 3 reflect the surrounding three quarters. Coefficients are normalized to be zero in period -2. The right panel shows the corresponding monthly series of the log-levels of taxable and tax-exempt spending generated by the continuous-time model in Section 5.

Figure 4 – Continuous-Time Model Dynamics around a Sales Tax Increase

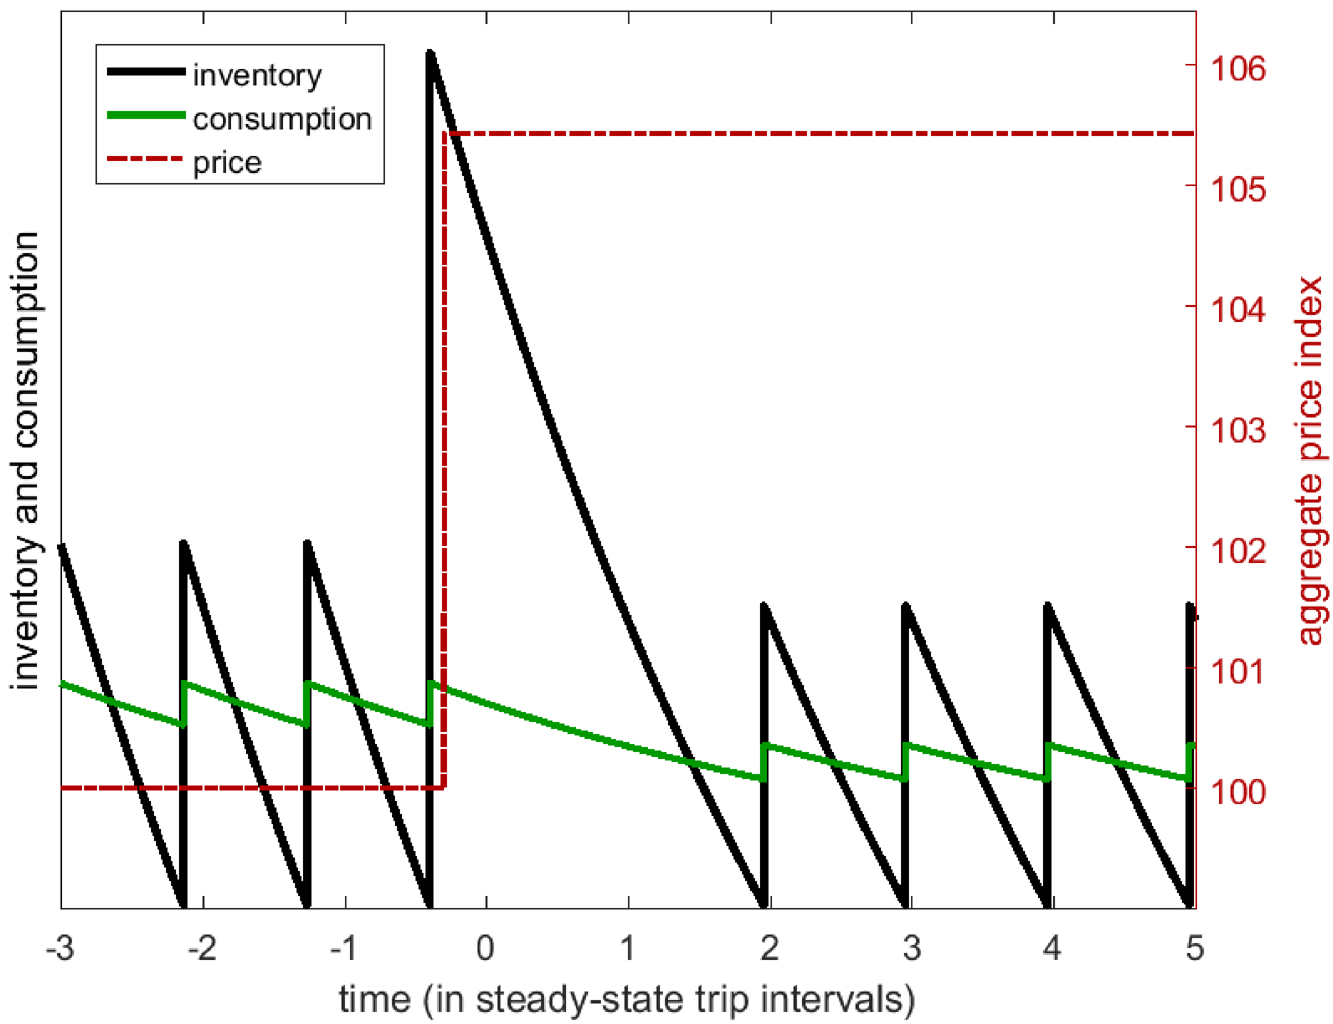

Notes: Figure plots the evolution of inventory $S(t)$ and consumption $C(t)$ of the composite good and the tax-inclusive aggregate price index $P(t)$ around a sales tax increase. To make this example as stark as possible, we use a very large tax change of 10% and an unrealistically large intertemporal substitution elasticity of $\sigma = 6$ and a low fixed cost of $\kappa = 2$. Otherwise, the model parameters are set as described in Section 5.

Table 1: Response of Taxable Spending and Shopping Frequency to a Sales Tax Increase

Data source:	Nielsen Consumer Panel (NCP)							
Dependent variable: log change of monthly taxable retail spending (Panels A and B).	A. Main Analysis			B. Robustness				
	Baseline	State tax rate only	Sales tax cuts	Household charac.	Business cycle	Drop Great Recession	State-period FE	No online spending
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δln(1 + total sales tax rate)	-2.036***		-1.719*	-2.034***	-2.082***	-2.012**	-2.269***	-2.224***
	(0.648)		(0.965)	(0.648)	(0.648)	(0.882)	(0.701)	(0.691)
Δln(1 + state sales tax rate)		-2.185**						
		(1.031)						
Period FE	Yes	Yes	Yes	Yes	Yes	Yes		Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household characteristics				Yes	Yes	Yes	Yes	Yes
Local unemployment rate					Yes	Yes	Yes	Yes
State-period FE							Yes	
Observations	4,137,927	5,928,468	4,114,413	4,137,927	4,137,927	3,285,747	4,137,886	3,565,145
R-squared	0.014	0.013	0.014	0.014	0.014	0.015	0.015	0.020

Data source:		Nielsen Retailer Panel (NRP)			Nielsen Consumer Panel (NCP)			
Dependent variables: log change of monthly taxable retail sales (Panel C) and the number of monthly store visits (Panle D).	B. cont.	C. Store Sales			D. Shopping Frequency (Log # of Trips)			
	No border shopping	Baseline	Business cycle	State-period FE	Baseline	Household charac.	Business cycle	Levels (# of trips)
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Δln(1 + total sales tax rate)	-2.145***	-2.814**	-2.794**	-3.265**	-1.479***	-1.479***	-1.455***	-20.164***
	(0.756)	(1.368)	(1.368)	(1.440)	(0.449)	(0.449)	(0.449)	(5.562)
Period FE	Yes	Yes	Yes		Yes	Yes	Yes	Yes
Household FE	Yes				Yes	Yes	Yes	Yes
Household characteristics	Yes					Yes	Yes	Yes
Local unemployment rate	Yes		Yes	Yes			Yes	Yes
State-period FE				Yes				
Store FE		Yes	Yes	Yes				
Observations	2,926,272	2,461,491	2,461,491	2,461,491	4,137,927	4,137,927	4,137,927	4,137,927
R-squared	0.024	0.140	0.140	0.150	0.020	0.020	0.020	0.026

Notes: Total sales tax rates combine all sales taxes within a ZIP code, including state, county, city, and special districts. The dependent variable in Panels A to C is monthly changes in logged household taxable spending as measured by Nielsen Consumer Panel and the Nielsen Retailer Panel data and the (level or log) of the number of shopping trips per month in Panel D. 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. The sample average number of trips per month in Column 16 is 15.1. Regressions span 2004-2014 for state sales tax rate changes (Column 2) and 2008-2014 for total 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 2: Storability and Intertemporal Substitution

Dependent variable: log change of monthly taxable retail spending by product group and state	A. Average Effect		B. Response by Storability			C. Poduct Groups by Purchase Frequency	
	Baseline	Lead	Binary	Linear	Quartiles	Non-storability: avg. # of purchases/month	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\Delta\ln(1 + \text{sales tax rate})$	-2.549**	-2.564**	-1.206	-1.538	-1.476	CARBONATED BEVERAGES	1.942
	(1.091)	(1.111)	(1.257)	(1.240)	(1.324)	CANDY	1.700
$\Delta\ln(1 + \text{sales tax rate})$, lead		0.834	0.092	-0.010	0.015	PAPER PRODUCTS	1.615
		(0.790)	(0.812)	(0.649)	(0.658)	PET FOOD	1.382
						ICE CREAM, NOVELTIES	0.967
$\Delta\ln(1 + \text{sales tax rate}) \times$ I(durable)			-2.979			SOFT DRINKS, NON-CARB.	0.906
			(2.915)			DETERGENTS	0.726
$\Delta\ln(1 + \text{sales tax rate}) \times$ I(durable), lead			1.629*			WRAPPING MATERIALS	0.612
			(0.927)			STATIONERY, SCHOOL SUP.	0.604
						ORAL HYGIENE	0.602
$\Delta\ln(1 + \text{sales tax rate}) \times$ Storability				-2.209		PET CARE	0.591
				(1.792)		HAIR CARE	0.558
$\Delta\ln(1 + \text{sales tax rate}) \times$ Storability, lead				1.802*		LAUNDRY SUPPLIES	0.516
				(1.037)			
						...	...
$\Delta\ln(1 + \text{sales tax rate}) \times$ Storability							
- Quartile 2					1.652*	BABY NEEDS	0.147
					(0.916)	ICE	0.146
- Quartile 3					-5.543*	FLORAL, GARDENING	0.143
					(2.911)	CHARCOAL, LOGS	0.110
- Quartile 4					-12.678**	COOKWARE	0.106
					(6.272)	FEMININE HYGIENE	0.099
$\Delta\ln(1 + \text{sales tax rate}) \times$ Storability, lead						FRAGRANCES - WOMEN	0.081
- Quartile 2					-0.300	PHOTOGRAPHIC SUPPLIES	0.076
					(0.489)	MEN'S TOILETRIES	0.075
- Quartile 3					1.951*	CANNING, FREEZING SUP.	0.069
					(1.127)	TOYS & SPORTING GOODS	0.065
- Quartile 4					14.910*	GRT CARDS/PARTY NEEDS	0.049
					(8.885)	SEWING NOTIONS	0.044
						SEASONAL	0.042
Period FE	Yes	Yes	Yes	Yes	Yes	SHOE CARE	0.041
Product FE	Yes	Yes	Yes	Yes	Yes		
State FE	Yes	Yes	Yes	Yes	Yes		
Observations	307,520	307,520	307,520	307,520	307,520	Sample mean (weighted)	0.846
R-squared	0.064	0.064	0.064	0.064	0.064	Sample standard deviation	0.602

Notes: The dependent variable is the monthly log change in taxable retail spending by product group and state across all households in the Nielsen Consumer Panel, 2004-2014. The main independent variable is changes in state sales tax rates. We drop ‘‘magnet data”, leaving 53 unique product groups. All regressions are estimated using least squares weighted by average sales per product group. Robust standard errors in parentheses are clustered at the state level.

Table 3: Tax-Exempt Spending Response to a Sales Tax Increase

Data source:	Nielsen Consumer Panel (NCP)					Nielsen Retailer Panel (NRP)		
Dependent variables: log change of monthly exempt retail spending (Panels A and B) and exempt retail sales (Panle C).	A. Main Analysis		B. Robustness			C. Store Sales		
	Baseline	State tax rate only	Household charac.	Business cycle	State-period FE	Baseline	Business cycle	State-period FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta\ln(1 + \text{total sales tax rate})$	-1.395***		-1.393***	-1.329***	-1.215**	-4.243***	-4.162***	-4.120***
	(0.513)		(0.513)	(0.513)	(0.557)	(1.249)	(1.248)	(1.363)
$\Delta\ln(1 + \text{state sales tax rate})$		-1.618**						
		(0.656)						
Period FE	Yes	Yes	Yes	Yes		Yes	Yes	
Household FE	Yes	Yes	Yes	Yes	Yes			
Household characteristics			Yes	Yes	Yes			
Local unemployment rate				Yes	Yes		Yes	Yes
State-period FE					Yes			Yes
Store FE						Yes	Yes	Yes
Observations	4,095,406	5,865,177	4,095,406	4,095,406	4,095,406	2,303,059	2,303,059	2,303,059
R-squared	0.015	0.014	0.015	0.015	0.016	0.141	0.141	0.150

Notes: See the description in Table 2.

Table 4: Evidence of Shopping Complementarity

	A. Revealed Cost Approach				B. Trip Complementarity		C. Placebo Test	
	frequent shoppers		infrequent shoppers		combined trips	separate trips	online purchases	
Dependent variable:	Δln(exempt)	Δln(taxable)	Δln(exempt)	Δln(taxable)	Δln(exempt)	Δln(exempt)	Δln(exempt)	Δln(taxable)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δln(1 + sales tax rate)	-0.010	-2.202**	-2.236*	-2.406*	-2.109**	-0.813	0.145	1.592**
	(0.756)	(0.910)	(1.191)	(1.451)	(0.828)	(1.16)	(0.464)	(0.709)
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,086,921	1,091,667	934,657	951,890	1,049,599	895,365	6,868,924	6,868,924
R-squared	0.016	0.017	0.022	0.020	0.017	0.018	0.004	0.005

Notes: Sales tax rates in Columns 1 to 6 combine all sales taxes within a ZIP code, including state, county, city, and special districts, while Columns 7 and 8 use state sales tax rates. The dependent variable is the monthly log-change of exempt or taxable retail spending as measured by Nielsen Consumer Panel data. Frequent shoppers in Columns 1 and 2 are consumers with average monthly trips above the 75th percentile (19 trips), while infrequent shoppers in Columns 3 and 4 have average monthly trips below the 25th percentile (9 trips). Columns 5 and 6 split the sample into households who possess the most disparate (Column 6) and most combined (Column 5) shopping trips in terms of taxable and exempt purchases. 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 5: Persistent Tax Incentives

	A. Fraction Spent in Alternative Tax Jurisdiction						B. Online Spending	
	short-run response (one-month change)				long-run (12-month change)		short-run	long-run
Dependent variable:	Δln(frac. alt. ZIP)		Δln(frac. in alt. state)		Δln(alt. ZIP)	Δln(alt. state)	Δln(online spending)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δln(1 + total sales tax rate)	-0.075	-0.182***			-0.674***			
	(0.072)	(0.056)			(0.093)			
Δln(1 + total sales tax rate) $\times$ avg. fraction in alt. ZIP3		1.497			5.484***			
		(0.951)			(1.507)			
Δln(1 + state sales tax rate)			-0.003	-0.017		-0.199***	1.703**	1.591**
			(0.027)	(0.019)		(0.034)	(0.824)	(0.791)
Δln(1 + state sales tax rate) $\times$ avg. fraction in alt. state				1.334		4.731		
				(2.639)		(3.907)		
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,231,065	4,231,065	6,049,454	6,049,454	2,510,373	4,010,860	6,868,924	3,010,794
R-squared	0.005	0.005	0.005	0.005	0.119	0.101	0.005	0.044
Average of interaction variable		0.079		0.017	0.079	0.016		

Notes: Regressions span 2004-2014 for state sales tax rate changes and 2008-2014 for total sales tax rate changes. Panel A examines changes in the monthly fraction of a household's retail spending in an alternative tax jurisdiction outside the household's residential 3-digit ZIP code, either an alternative 3-digit ZIP code or an alternative state. The dependent variable in Columns 1 to 4 are monthly changes while Columns 5 and 6 use 12-moth changes (e.g., change from March 2013 to March 2014). Panel B examines changes in the log of total online spending, including mail orders. 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 6: Salience and Announcement Effects

	A. Salience Effects						B. Announcement Effects	
	newspaper coverage			ballot-induced tax changes				
Dependent variable:	Δln(total)	Δln(taxable)	Δln(exempt)	Δln(total)	Δln(taxable)	Δln(exempt)	Δln(total)	Δln(total)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δln(1 + sales tax rate)	-1.738***	-2.124**	-1.572**	-1.526**	-2.238*	-1.310**		
	(0.581)	(1.053)	(0.603)	(0.687)	(1.179)	(0.591)		
Δln(1 + sales tax rate) × Score(newspaper coverage)	-0.361***	-0.336	-0.439**					
	(0.110)	(0.257)	(0.166)					
Δln(1 + sales tax rate) × I(state ballot proposition)				-4.195***	-4.765**	-5.043***		
				(1.050)	(2.038)	(0.889)		
I(date tax rate change proposed)							-0.529	-1.706
							(0.330)	(1.444)
I(date tax rate change proposed) × I(ballot proposition failed)								1.434
								(1.493)
Score(newspaper coverage of state sales tax changes)	-0.001	-0.001*	0.001					
	(0.001)	(0.001)	(0.001)					
I(date ballot proposition failed)				0.022***	0.030***	0.022***		
				(0.005)	(0.009)	(0.005)		
I(ballot proposition failed)								-0.002
								(0.006)
Period FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,822,806	5,777,878	5,865,177	5,865,949	5,928,421	5,777,966	5,860,476	5,860,476
R-squared	0.016	0.015	0.014	0.014	0.012	0.013	0.014	0.014

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 state sales tax rates with the level of newspaper coverage (measured as the demeaned ratio of articles mentioning sales taxes to the total number of articles in newspapers within the state covered by Access World News, normalized by its standard deviation). 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.