

Raising the Stakes: Experimental Evidence on the Endogeneity of Taxpayer Mistakes

Naomi Feldman Jacob Goldin Tatiana Homonoff*

May 24, 2017

Abstract

Recent evidence suggests consumers fail to account for taxes that are excluded from a good’s displayed price. What is less understood is whether and how such “salience effects” depend on the magnitude of the tax. We conduct a laboratory shopping experiment with real stakes to study the effect of tax size on salience. We find no evidence that salience effects decline as the tax rate increases; we document a statistically significant salience effect at a tax rate that is considerably larger than the tax rates at which such effects have been previously documented. In fact, our results are more consistent with the hypothesis that higher taxes make consumers less attentive (at least for the range of taxes we consider). This result can be explained by a confirmation bias theory of salience: consumers tend to disregard information (like a tax) that does not align with their intention to purchase an item, and this lack of alignment increases in the size of the tax.

Introduction

A growing body of evidence suggests that the presentation of a tax affects taxpayer decision-making. With respect to commodity taxes, several studies suggest that consumers fail to fully account for taxes that are not included in the displayed price of the taxed good, and, as a result, choose to

*Feldman: Federal Reserve Board. Goldin: Stanford Law School. Homonoff: NYU, Wagner School of Public Service. We thank Henry Farber, David Lee, Alexandre Mas, Baxter Oliphant, Daniel Reck, Bradley Ruffle, Erez Semaria, and numerous seminar and conference participants for helpful suggestions. Financial support from the Princeton Industrial Relations Section and Princeton Lab for Experimental Social Science (PLESS) is gratefully acknowledged. The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Reserve Board. Any remaining errors are our own.

purchase more of the taxed good than when the tax is included in the displayed price. These so-called “salience” effects have been documented for commodity taxes both experimentally and in observational studies.¹

In this paper, we move beyond the question of whether individuals behave differently when faced with high versus low salience taxes and instead investigate a potential determinant for whether taxpayers account for low salience taxes when making purchasing decisions: the magnitude of the tax being levied.² Understanding the relationship between salience and tax size is important for two reasons, one positive and one normative. First, it can help differentiate among plausible positive models of decision-making, such as those in which decision-makers trade-off the utility costs and benefits of alternate decision-making strategies. For example, when the utility benefits of accounting for low salience taxes outweigh the utility costs of doing so (e.g., cognitive costs, opportunity costs of diverted attention), consumers may become more attentive when facing higher taxes. Alternatively, decision-making biases may cause individuals to become less attentive when facing higher taxes if these taxes are viewed as unpleasant and at odds with the decision-maker’s intention to purchase the taxed good. This possibility is based upon the widespread phenomenon known in social psychology as “confirmation bias.” Well known manifestations of confirmation bias include ignoring information that contradicts one’s political views as well as the overweighting of first impressions.³

The second reason to care about the relationship between tax size and salience is a normative one. Several recent papers find that low salience tax instruments can offer efficiency and distributional benefits in certain circumstances (Goldin, 2015; Goldin and Homonoff, 2013). However, the potential benefits of such taxes are limited by the extent to which consumers become more or less attentive to them as governments levy them at higher rates.⁴ Thus, understanding how salience varies based on the tax rate is crucial for assessing the benefits of employing such taxes to further policy goals.

¹For experimental evidence, see Chetty, Looney and Kroft (2009); Feldman and Ruffle (2015); Hayashi, Nakamura and Gamage (2013). Observational studies have studied commodity tax salience in the context of purchases of beer (Chetty, Looney and Kroft, 2009), hybrid cars (Gallagher and Muehlegger, 2011), soft drinks Colantuoni and Rojas (2015), cigarettes (Goldin and Homonoff, 2013), and airline tickets (Bradley and Feldman, 2015). Outside of the commodity tax context, salience effects have been documented for highway tolls (Finkelstein, 2009), property tax payments Hayashi (2014); Bradley (2017), and shipping and handling fees (Brown, Hossain and Morgan, 2010).

²A “high salience” tax is a tax for which the after-tax price is prominent, such as a VAT or excise tax, where the good’s displayed price includes the tax. In contrast, a “low salience” tax is one for which the full after-tax price is not prominent, such as a sales tax, where the tax is added to the displayed price at the conclusion of the purchase.

³See Nickerson (1998) and Hart et al. (2009) for comprehensive reviews of the confirmation bias literature.

⁴In fact, if there exists some sufficiently large tax rate at which all consumers become attentive, Reck (2015) shows that there must exist some rate for the low salience tax above which further increases generate more deadweight loss than a high salience tax.

In extreme cases, it seems almost certain that individuals would account for a low salience tax when making purchasing decisions. For example, if a \$1,000 per-unit tax were imposed on one brand of toothpaste, it is likely that almost everyone would take the tax into account when deciding which brand of toothpaste to buy. The more interesting question is whether taxpayers become attentive to the tax at high, yet policy-relevant rates even when the “stakes” are relatively low, as is the case in most everyday shopping for food, toiletries, basic household items and (some) apparel.⁵

A small but growing literature that is closely related to ours examines whether behavioral biases fade as the stakes involved in the decision grow larger. A classic study in this area is Camerer and Hogarth (1999), which provides evidence that increasing the financial stakes in experiments does not induce subjects to avoid cognitive biases or other errors. More recently, Haggag and Paci (2014) study how changing the default tip suggestions in taxicabs affects tipping behavior; they find that larger defaults result in higher tips among those who choose to tip but also a higher number of customers opting out. Similarly, Brown et al. (2013) documents similar behavior with respect to decisions about setting the temperature; when nudged towards slightly cooler temperatures, people do not adjust the thermostat but when nudged towards much cooler temperatures then people change the setting.⁶ Within the domain of taxation, Chetty et al. (2011) find that taxpayers respond more to income tax incentives when the utility stakes of ignoring the incentives are greater. To our knowledge, we are the first to investigate such phenomena in the context of commodity taxation.⁷

To investigate the relationship between salience and the size of a tax, we designed a lab experiment with real stakes to simulate an online shopping environment. Subjects were provided with \$18 and an opportunity to purchase various products. Depending on the treatment group to which

⁵Our results are unlikely to speak to decision-making about large, durable purchases, such as computers or automobiles.

⁶A closely related literature looks to see if people’s preferences become less altruistic as the amount at stake is increased. For example, Camerer (2003) finds that people’s behavior in ultimatum games is largely unaffected by changing the stakes.

⁷An informative new paper by Taubinsky and Rees-Jones (2016) also reports results from an experiment designed to test how salience varies based on the tax rate. Specifically, the paper solicits from individuals information about the maximum pre-tax price they would be willing to pay for a particular product, varying the sales tax rate within and between individuals to identify changes in salience. In contrast to us, they find that tax salience declines as the tax rate increases. This paper builds on ours by utilizing a larger and more representative sample of the United States population and by developing a theoretical framework for drawing welfare conclusions from the experimental findings. On the other hand, the decision problem they pose to survey-takers may be less likely to generalize to real-world purchasing settings. For example, the extent to which one considers a tax when contemplating a reservation purchase price may differ from how one considers a tax when making the concrete decision of whether or not to purchase an item.

a subject was assigned, the subject made purchasing decisions presented with tax-inclusive (TI) or tax-exclusive (TE) prices, at high (22%) or low (8%) tax rates.⁸ To measure salience effects at a given tax rate, we compared the gap between subjects' expenditures at TE versus TI prices. By observing this quantity at the low tax rate and the high tax rate, we are able to measure changes in the salience of the tax across treatments. In order to isolate variation in behavior caused by changes in the cost of accounting for the tax, we adjusted the pre-tax price of the available products so that after-tax prices were equalized between the high- and low-tax treatment groups. We interpret the results of the experiment using a simple model to identify whether observed differences in consumption between treatment groups are due to changes in tax size or salience.

Our results suggest that economically significant salience effects persist even at the high (22%) tax rate we consider. At that rate, the gap in expenditures between TE and TI prices is positive (approximately \$1.10) and statistically significant ($p < 0.01$). Notably, the estimated expenditure gap for the low (8%) tax rate, while also positive (approximately \$0.20), is smaller than at the high tax rate ($p < 0.01$) and statistically indistinguishable from zero. Drawing on our theoretical results, we isolate the portion of the observed expenditure gap change that is due to changing tax salience and find that the higher tax is associated with a smaller degree of salience, although this difference is not statistically significant. This pattern of estimates provides no support for the hypothesis that tax salience increases in the size of the tax for the range of taxes we consider, and suggestive evidence in favor of models of tax salience based on confirmation bias.

The remainder of the paper proceeds as follows. Section I develops the theoretical framework. Section II describes the experimental design. Section III presents descriptive statistics concerning the experiment subjects. Section IV presents the experimental results. Section V concludes.

I. Theoretical Framework

This section develops a simple model to motivate the experimental design and to facilitate the interpretation of the experimental results. A representative consumer chooses between two goods, x and y . Good x has displayed price p and is subject to an additional specific tax t . Good y is

⁸We choose 8 percent as the low tax rate because this is a typical sales tax rate imposed on the East Coast of the U.S., where the experiment was conducted. The 22% high tax is more typical of the VAT rates imposed in Europe.

untaxed and we normalize its price to 1. The budget constraint is given by

$$(p + t)x + y = Z \tag{1}$$

where Z represents the consumer’s (fixed) endowment. The consumer’s demand for x and y is a function of the displayed price and the tax: $x = x(p, t)$ and $y = y(p, t)$.

As in Chetty, Looney and Kroft (2009), the salience of a tax, θ_t , is defined in terms of how the tax affects consumer demand for the taxed good:

$$\theta_t = \frac{\frac{\partial x}{\partial t}}{\frac{\partial x}{\partial p}} \tag{2}$$

A fully-salient tax ($\theta_t = 1$) is one to which consumers respond as if the change had been to the pre-tax price; when $\theta_t = 0$, a change in the tax does not cause any change in demand for the taxed good. We index θ by t to make explicit the possibility that the size of a tax affects the salience associated with it.

As discussed above, the sign of $\frac{\partial \theta_t}{\partial t}$ is theoretically ambiguous. Under a bounded rationality model of decision-making, higher taxes would increase θ because the utility stakes of neglecting the tax would be greater than when the tax was small.⁹ Alternatively, models of confirmation bias might predict that the unpleasantness of accounting for a tax on an intended purchase would grow as the size of the tax increases, causing consumers to become less attentive as the tax grows in size. Finally, salience effects may simply reflect innate decision-making biases or heuristics that do not vary based on the tax size, at least under the range of tax rates we observe.

To investigate differences in salience across tax sizes, we experimentally generated consumption data across variation in tax sizes (high versus low) and tax designs (tax-inclusive versus tax-exclusive prices). To isolate differences in consumption arising from differences in the utility cost of accounting for the tax, we held the after-tax price constant across treatments. Let t_j denote the low or high tax, $j \in \{l, h\}$, with $t_h > t_l$. Let p_j denote the displayed price corresponding to t_j for $j \in \{l, h\}$. By design, p_h is chosen so that the after-tax price is constant at the high and low taxes: $t_h + p_h = t_l + p_l$. Under the tax-exclusive (TE) price presentation, consumption of x is given by $x(p_h, t_h)$ under the high tax treatment and by $x(p_l, t_l)$ under the low tax treatment. In contrast,

⁹See Chetty, Looney and Kroft (2007) for a formal model along these lines.

when prices are tax-inclusive (TI), the displayed price includes the tax and consumption is given by $x(p_h + t_h, 0) = x(p_l + t_l, 0)$.

Our primary quantity of interest will be the *expenditure gap*, $g(t)$, which refers to the difference between the amount expended on x at TE versus TI prices at tax t :

$$g(t) \equiv (p + t) x(p, t) - (p + t) x(p + t, 0) \quad (3)$$

Adding and subtracting $(p + t) x(p, 0)$ and taking Taylor approximations allows us to approximate (3) as $g(t) \approx (p + t) t \left(\frac{\partial x}{\partial t} - \frac{\partial x}{\partial p} \right)$.¹⁰ Substituting (2) into this equation allows us to express the expenditure gap as a function of salience:

$$g(t) \approx -(1 - \theta_t) (p + t) t \frac{\partial x}{\partial p} \quad (4)$$

Equation (4) highlights that differences in the observed expenditure gap at the high versus low taxes do not necessarily reflect changing levels of salience; a constant level of salience would mechanically yield a larger expenditure gap at the high tax, as the effect on consumption of ignoring a large tax is greater than the effect of ignoring a small tax. However, we can scale $g(t)$ by the size of the tax and the after-tax price to back out changes in salience: $\frac{g(t)}{t(p+t)} = -(1 - \theta_t) \frac{\partial x}{\partial p}$. Using the fact that the after-tax price is the same under the high- and low-tax treatments, we obtain:

$$UR \equiv \frac{1 - \theta_{t_l}}{1 - \theta_{t_h}} \approx \frac{g(t_l)}{g(t_h)} \frac{t_h}{t_l} \quad (5)$$

This quantity, $\frac{1 - \theta_{t_l}}{1 - \theta_{t_h}}$, can be interpreted as the *Unresponsiveness Ratio* (UR) between the two taxes. When both taxes have the same salience, i.e., when consumers are equally unresponsive to each, we have $UR = 1$. In contrast, when the salience of a tax is increasing in the tax's size, this implies $UR > 1$.¹¹

¹⁰Note that $\frac{\partial x}{\partial p}$ and $\frac{\partial x}{\partial t}$ are evaluated at $(p, 0)$, so that comparing g at different taxes (with different displayed prices) requires assuming local linearity in demand for x . If instead it turned out to be the case that $\frac{\frac{\partial x}{\partial p}|_{p+t_h}}{\frac{\partial x}{\partial p}|_{p+t_l}} > 1$, our estimate of the change in salience would be biased upwards (i.e. our estimated unresponsiveness ratio, defined below, would be greater than the true unresponsiveness ratio). This might occur, for example, if subjects ignored the tax and determined the perceived 22% discount beyond the base price to be a substantially better deal than the perceived 8% discount beyond the base price.

¹¹Note that $UR \rightarrow \infty$ if a low salience tax is increased up to the point at which all taxpayers become perfectly responsive.

Finally, to make the experiment more realistic, we presented taxes as ad valorem rates rather than specific taxes. Let τ_l indicate the low tax rate and τ_h indicate the high tax rate, where $\tau_l = t_l/p_l$ and $\tau_h = t_h/p_h$. Because the after-tax prices in our experiment were chosen to be equal, we know

$$(1 + \tau_l)p_l = (1 + \tau_h)p_h \tag{6}$$

Substituting this identity into (5) yields

$$UR = \frac{g(\tau_l)}{g(\tau_h)} \frac{\tau_h (1 + \tau_l)}{\tau_l (1 + \tau_h)} \tag{7}$$

We will utilize equation (7) to interpret the results of our experimental analysis.

II. Experimental Design

A. Set Up

The experiment consisted of 16 total rounds of purchasing decisions by experimental subjects (undergraduate students at Princeton University). In each round, subjects were endowed with \$18 and presented with consumption goods that they could purchase. Subjects were informed that they could purchase as many of each good as they wished (up to their \$18 endowment) and that they would get to keep however much of their endowment was left over. At the end of the experiment, one round was randomly chosen for payment. The actual goods and any leftover money from the \$18 were distributed in the lab. These aspects of the design were similar to Feldman and Ruffle (2015).

As with all laboratory experiments, an important concern is with external validity. To make the experiment as realistic as possible, we attempted to simulate a typical online shopping experience. We purchased an internet domain and directed subjects to proceed to the website, which is where they made their purchasing decisions. The sequence of events was as follows: first, subjects were presented with an instruction screen (Figure 1). The instruction screen stated that all goods were subject to a tax and stated the tax rate (either 8 percent or 22 percent). Additionally, subjects were informed either that the displayed prices already included the sales tax (the TI treatment) or did not include the sales tax (the TE treatment). In the latter case, the instructions stated that

the sales tax would be added upon the final purchase of the good.

After viewing the instructions, subjects proceeded to the first round. The main portion of the experiment consisted of two parts, with four rounds per part. Each round began with a shopping stage, in which subjects were presented with a set of products and prices available for purchase (Figure 2).¹² The menu of available consumption goods varied by round. Of the ten total products available for purchase, a menu of five were available for purchase in a single round. There were four such menus that were available.¹³

After selecting how many of each good to purchase, subjects proceeded to the checkout screen, which summarized the selected purchase (Figure 3). The checkout screen always contained the after-tax prices, regardless of whether the prices on the previous screen were TE or TI prices. At checkout, subjects had the option of either finalizing their purchase or going back to modify it. Additionally, for subjects who attempted to spend more than their \$18 endowment, the checkout screen informed them that they were required to go back and select fewer goods. Once the purchase was finalized, subjects proceeded to the shopping stage for the next round. At the conclusion of the experiment, one of the rounds was randomly selected and subjects received their purchased products and leftover endowment from that round.^{14,15}

B. Experimental Groups

Before the experiment began, subjects were assigned to one of four experimental treatment groups: (1) high-tax treatment, (2) high-tax control, (3) low-tax treatment, and (4) low-tax control. Those in the high-tax group faced a tax rate of 22 percent throughout the experiment whereas the tax rate was 8 percent for those in the low-tax group.¹⁶ It is important to note that while the low tax

¹²To minimize the corner solution in which subjects opted to keep their entire endowments, we offered the products at pre-tax prices that were roughly 50 percent of their normal retail pre-tax prices.

¹³The menus were: Menu A: {Nutella, Coconut Water, Clif Bar, Pens, Laundry Detergent} Menu B: {Starbuck's Via Instant Coffee, Trail Mix, Ferrero Rocher Chocolates, Coconut Water, Laundry Detergent} Menu C: {Starbuck's Via Instant Coffee, Clif Bar, 5-hour, Pringles, Nutella} and Menu D: {Trail Mix, Ferrero Rocher Chocolates, 5-hour, Pringles, Pens}.

¹⁴To avoid imposing transaction costs associated with acquiring the goods (e.g., traveling to a store, exchanging a voucher for the goods), we purchased all of the goods ahead of time and paid subjects according to their choices at the conclusion of each experimental session.

¹⁵Note that in our experimental set-up, cash is the untaxed good described in the theoretical model in Section I.

¹⁶We chose not to vary the tax rate within subject over the course of the experiment to avoid calling attention to the tax rate and artificially increasing its salience.

rate is within the range of sales tax rates levied in the US, the high tax rate is not.¹⁷ Sales tax (i.e. VAT) rates observed outside of the US are typically much higher; for example, the average VAT rate in the European Union was 22 percent at the time of our study. Therefore, while our high tax rate is likely larger than rates previously experienced by our participants, it is not outside the policy-relevant range.¹⁸

The distinction between the treatment and control groups is determined by the rounds in which the displayed prices were tax inclusive or tax exclusive. When prices were tax-inclusive, the displayed prices included the tax at both the shopping and checkout stages. When prices were tax-exclusive, the displayed prices were pre-tax; the taxes did not get added until they proceeded to the checkout. The experiment was divided into two parts, with half of the rounds in each part.¹⁹ For the treatment groups, the displayed prices were TE in Part I and TI in Part II. For the control groups, the displayed prices were TI in both Parts I and II.²⁰ Each subject received each product menu twice during the main portion of the experiment: once in Part I and once in Part II (the order in which subjects received each menu in each set of four rounds was randomized to reduce any order effect). Between Parts I and II, subjects in both treatment groups were informed that they would be shopping in a new store for the remainder of the experiment, and that prices in the new store included the sales tax (Figure 4).

C. Product Prices

Table 1 presents the pre- and post-tax prices for the experimental goods. The particular goods were chosen to appeal to a student population based on products available at the University store, and prices were chosen to facilitate comparison between the various treatment groups. In particular, during the main portion of the experiment, each treatment group faced the same after-tax price of each good. For example, the after-tax price for Nutella was \$3.49. Treatment subjects assigned to

¹⁷There are many localities, for example, Seattle, Oakland or Chicago that levy sales taxes (combined state and local rates) at levels approaching 10 percent.

¹⁸Our experiment was run in New Jersey which had a sales tax rate of 7 percent at the time of our study. Additionally, New Jersey exempts food purchases from the sales tax. We chose not to add this level of complexity to our experimental tax rates and were explicit in the instructions that all products were subject to the tax.

¹⁹To minimize the repetitiveness of the experiment, halfway through the experiment all subjects were shown two short videos containing advertisements for laptop computers and asked which advertisement they preferred.

²⁰No treatment group went in the reverse order because we expected that altering the salience of the tax between rounds would increase its salience; going from TE to TI avoids that problem as long as the displayed prices in the final rounds are already fully salient.

the 8 percent tax group had the pre-tax price of \$3.23 displayed to them while shopping in Part I and the after-tax price of \$3.49 displayed to them while shopping in Part II. Similarly, treatment subjects assigned to the 22 percent tax saw the pre-tax price of \$2.86 while shopping in Part I and the after-tax price of \$3.49 while shopping in Part II. In contrast, subjects assigned to the control groups saw the after-tax price of \$3.49 while shopping throughout the entire experiment. Because the after-tax price was the same for each treatment group, all subjects saw the same \$3.49 price during checkout for all rounds.

Finally, we included an additional four rounds in each Part at an increased price level for the purpose of testing whether demand was downward sloping in after-tax price changes (see Section IV.B.2). The prices in the high price rounds are displayed in Table 2.

III. Experimental Data

A. Experimental Subject Characteristics

The experimental subjects were 227 Princeton University undergraduate students. Table 3 presents demographic information for the sample. On average, subjects were 20 years old and 70 percent of them were in either their junior or senior year of college. Thirty percent of subjects were male, half were minorities, and half described their financial status as upper class or wealthy. These figures are roughly similar to the school average, though our sample contained a higher proportion of upperclassmen and females than the general student body.

B. Randomization Checks

Subjects were randomly assigned into one of four treatment groups. To maximize statistical power, one-third of subjects were assigned to each of the two treatment groups and one-sixth of subjects were assigned to each of the two control groups. Although the realized distribution of subjects across treatment groups did not precisely coincide with these assignment probabilities,²¹ Table 3 shows that subjects' demographic characteristics did not appear to systematically vary with their

²¹This likely occurred because the treatment group was assigned with replacement and the assignment probabilities were held constant over the course of the experiment.

treatment status, consistent with the randomization working properly.²²

C. Purchasing Behavior

Table 4 summarizes the purchasing decisions of the experimental subjects. On average, subjects spent \$6.16 per round, just over one-third of their \$18 budget. Only 13 percent of subjects chose not to purchase any products throughout the experiment. On average, subjects purchased at least one item in 67 percent of their rounds. Forty-three percent of subjects purchased at least one item in every round. The distribution of expenditures per round is displayed in Figure 5.

One concern is that if subjects would prefer to spend more than \$18 on product purchases, that would mechanically limit salience effects as they would be unable to increase the amount of their purchases even if they failed to account for the after-tax price. However, subjects spent over \$15 in only 17 percent of rounds on average, and only 7 percent of subjects did so in every one of their rounds. This suggests that most subjects were not at the corner solution of allocating their maximum allowable budget to the available products.

IV. Empirical Results

A. Tax Salience Effects

This section turns to our primary question: how do subjects respond differently to tax-inclusive (TI) versus tax-exclusive (TE) prices, and does the effect vary by tax rate? To investigate this, we first estimate the expenditure gap (the difference in expenditures between TE and TI prices) for subjects in each tax group. Next, we use the theoretical results from Section I to translate the estimated expenditure gap into the Unresponsiveness Ratio, which measures changes in salience between the high and low tax rates.

1. Expenditure Gap

Because subjects in the treatment group were presented with TE prices during Part I and TI prices during Part II, we can use a within-subject design to estimate the effect of tax salience on purchasing

²²In particular, an F-test fails to reject equality across the four experimental groups for each of these demographic variables.

behavior. As subjects may have changed their purchasing behavior as the experiment progressed for reasons unrelated to differences in salience between the early and late rounds, we include round fixed effects to capture differences in purchasing behavior associated with the timing of the purchasing decision within the experiment. The round fixed effects are identified by the control groups, who do not experience variation in salience over the course of the experiment.

Table 5 presents the unadjusted mean expenditure per round for Part I and Part II by tax rate and treatment group. Comparing the means for these groups relies on between-subject comparisons, which, due to the significant heterogeneity in individual preferences, are not statistically distinguishable from one another. To utilize our experimentally generated within-subject data, Table 6 presents the results of a regression analysis. As a baseline, Column 1 pools the high- and low-tax groups to check for the presence of a salience effect. The econometric model in the pooled analysis is given by:

$$y_{imr} = \beta TE_{ir} + \alpha_{im} + \gamma_r + \varepsilon_{imr}$$

In this model, y_{imr} denotes total expenditures for subject i facing menu of goods m in round r , TE_{ir} equals one in rounds where taxes are excluded from the displayed price and zero in rounds where taxes are included in the price, α_{im} are subject-by-menu fixed effects, and γ_r are round fixed effects. In this model, the β coefficient is a within-subject measure of the difference in expenditures between rounds that include versus exclude taxes from the displayed price (the “expenditure gap”).

The results of this analysis are consistent with past findings from the salience literature. In particular, subjects spend approximately \$0.57 more when taxes are excluded from the displayed price, suggesting a salience effect.

Turning to the un-pooled analysis, the second column in Table 6 repeats this analysis but allows the salience effects to differ by tax rate. The econometric model for the un-pooled analysis is given by:

$$y_{imr} = \beta_{low} TE_{ir} * LowTax_i + \beta_{high} TE_{ir} * HighTax_i + \alpha_{im} + \gamma_r + \varepsilon_{imr}$$

where $LowTax_i$ is an indicator variable for being assigned to the 8 percent tax rate and $HighTax_i$ is an indicator variable for being assigned to the 22 percent tax rate.

Column 2 of Table 6 documents a striking disparity in sensitivity to tax exclusive pricing by tax rate. Both treatment groups spent more when presented with TE prices, but the magnitude of the difference was much greater for those in the high tax group. The estimated coefficient for β_{high} suggests that the high tax group spent approximately \$1.07 more per round when prices were TE than when prices were TI (the effect was statistically significant at the 1 percent level). In contrast, we estimate that the low tax group spent only \$0.17 more per round under TE prices, and the effect was not statistically distinguishable from zero. These results suggest that the expenditure gap increases with the tax rate. However, as discussed in Section I, the expenditure gap alone does not measure subjects' attentiveness to tax rates; the key parameter needed for answering that question is the Unresponsiveness Ratio.

2. Estimated Unresponsiveness Ratio

To determine whether salience varied by tax rate, we can substitute the estimated expenditure gap into the Unresponsiveness Ratio formula derived in Section I (Equation 7):

$$UR = \frac{g(\tau_l)}{g(\tau_h)} \frac{\tau_h(1 + \tau_l)}{\tau_l(1 + \tau_h)}$$

From the regression, we know the ratio of expenditure gaps: $\frac{g(\tau_l)}{g(\tau_h)} = \frac{0.17}{1.07} = 0.16$. Substituting the other parameters into the formula yields

$$\hat{UR} = (0.16) \frac{(0.22)(1.08)}{(0.08)(1.22)} = 0.40$$

with a standard error of 0.53.²³

This point estimate suggests that salience is *decreasing* in tax size, consistent with a confirmation bias. As previously mentioned, a confirmation bias hypothesis maintains that an individual generally likes to avoid personally disquieting information, and, the strength of a bias in the interpretation of evidence increases with the degree to which the evidence relates directly to a dispute in which one has a personal stake. In our context, the tax tends to undermine one's intentions to buy, that is, by increasing the final price, the tax provides a negative signal to the purchaser thereby causing the purchaser to disregard or downweight this evidence. Thus, this model would suggest that consumers

²³We calculate the standard error of this parameter using the delta method.

would be even more likely to disregard information and feedback when facing higher taxes (again, because these taxes go against the vested intent to purchase and this is all the more true as the size of the tax grows.). However, the estimate is not sufficiently precise to allow us to reject a value of the Unresponsiveness Ratio that is greater than or equal to 1 ($p \approx 0.25$).

To give some context for interpreting this result, if a commodity tax in the neighborhood of 8 percent was associated with a level of salience of 0.35 (as reported in Chetty, Looney and Kroft (2009)), our results suggest that increasing the tax rate by 175 percent – up to 22 percent – would result in a new level of salience of about 0.14. That is, at the 22 percent tax rate, consumers would need to face a tax change that is six times as large as a price change in order to induce similar behavioral responses.

B. Robustness Checks

1. Understanding the Experiment

The findings in Section IV.A are robust to several sample selection checks. One concern with all lab experiments is that subjects did not understand the instructions of the experiment. We included two checks on subjects' knowledge of the directions in a post-experiment questionnaire. The first check asked subjects about their perceptions of their own understanding of the experiment and the second check asked subjects to recall the tax rate that they faced in the experiment. If we drop the 27 subjects who reported less than perfect understanding of the directions or the three subjects who incorrectly reported their assigned tax rate, the main results are almost completely unchanged.

2. Price Responses

In addition to the eight experimental rounds that were used to estimate the expenditure gap and the Unresponsiveness Ratio, we also ran another eight rounds where we increased the prices of each of the goods by a fixed percentage in order to assess the effect of pre-tax price increases on subjects' purchasing decisions.²⁴ Relative prices among the goods remained unchanged. Following the law

²⁴Whereas the after-tax prices were held constant across treatment groups at the low price level, subjects in the high- and low-tax treatment groups faced different after-tax prices in the high-price level rounds (and therefore were not used in the expenditure gap calculations of Section IV.A). The high price was equal to 1.22 times the low price for the high tax group and 1.08 times the low price for the low tax group – these prices were chosen to make the price variation comparable to the observed tax variation.

of demand, we would expect that as long as subjects were paying attention to their purchasing decisions with which they were faced, overall demand would fall. Such a finding would provide further confidence that the subjects were otherwise acting rationally and thoughtfully.

Like the low-price rounds described earlier, each subject was exposed to a high after-tax price level for each product menu. This allows us to compare per-round expenditures within subject between the high and the low prices. The econometric model is as follows:

$$y_{imr} = \beta * HighPrice_{ir} + \alpha_{im} + \gamma_r + \varepsilon_{imr}$$

In this model, y_{imr} denotes total expenditures for subject i facing menu of goods m in round r , $HighPrice_{ir}$ equals one in rounds where subjects face high prices, α_{im} are subject-by-menu, and γ_r are round fixed effects. Thus the β coefficient is a within-subject measure of the difference in expenditures between rounds with high and low price levels.

Table 7 presents the results from this analysis. The results are broken up by treatment group because the difference between the high and low price levels was greater for the high tax groups (for whom the price difference was 22%) than for the low tax groups (for whom the price difference was 8%).

The results confirm that subjects' demand was downward-sloping: both groups purchased more in the low-price rounds than in the high-price rounds. As expected, the difference was greater for the high-tax groups, who faced a larger difference in prices.

The results in Table 7 also allow us to investigate potential non-linearities in subjects' change in demand in response to an 8% price increase and a 22% price increase. Mean expenditures in the low-price rounds with TI prices was \$5.86. Using the coefficients from Table 7, this implies that the change in demand associated with an 8 and 22 percent price increase represented (respectively) a 10.0 and 15.7 percent decline in expenditures, or own-price elasticities of 1.25 and 0.71. Although we cannot reject the null hypothesis that these two elasticities are in fact equal to one another, the point estimates suggest that subjects may also have been relatively more responsive to the perceived price change induced by the excluding the high tax from the posted price, as compared to the perceived price change associated with excluding the low tax from the posted price.

3. Censoring

As mentioned earlier, our results will be mechanically biased toward finding no evidence of tax salience if our subjects are censored from above or below. For example, subjects who are not interested in purchasing any of the products we offer will have the same expenditure under both TE and TI pricing. Similarly, subjects who prefer to buy as many products as their budget allows will also have the same expenditure under the two tax-framing treatments. Table 8 determines whether this censoring biased our salience estimates by excluding these two groups of subjects. Column 1 excludes the 25 subjects who never purchased any items in any round and column 2 drops the 12 subjects who spent over \$15 in every round. As expected, the salience estimates increase, though our results are largely unaltered under each of these tests.²⁵

Additionally, we analyze whether the presentation and size of the tax affected purchasing decisions separately by the extensive and intensive margins, that is, whether people purchased anything at all and how much those making purchases spent. The results are presented in Table 9. On both margins, the results are qualitatively similar as in the main analysis. On the extensive margin, subjects were more likely to make a purchase when facing TE prices, and this effect was approximately twice as large at the high tax rate than at the low tax rate (although the difference was not statistically significant). On the intensive margin, those subjects making purchases while facing the high tax rate spent approximately 16 cents more per round when facing TE prices, whereas the effect was close to zero and statistically insignificant for subjects at the low tax rate.

4. Tax Aversion

When interpreting our results, one concern is that subjects have different preferences for paying taxes than they do for other types of prices. One example of this would be if subjects were tax-averse, i.e., they are more likely to avoid taxes than other costs of the same amount. If this were the case, treatments that highlighted the size of the tax rather than the overall after-tax price may lead subjects to spend less on their purchases, independent of any salience effect. To test for the presence of tax aversion, we use a between-subjects design to compare the expenditure of subjects in the high- and low-tax groups. We include only tax-inclusive rounds; hence both the after-tax

²⁵Additional (unreported) results from censored least absolute deviation estimation also confirm our main results and estimate a slightly higher salience effect for the high tax rate.

and displayed prices were constant. All that varies in this analysis is how much of the after-tax price is tax as opposed to pre-tax price. If subjects are tax-averse, we would expect to observe that they spend more when the tax rate is 8 percent compared to 22 percent. Our results are presented in Table 10. Column 1 includes only control group subjects who were presented with TI prices throughout the entire experiment while column 2 extends the analysis to include the TI rounds of treatment subjects. The results show that the differences in expenditure are not significantly different when the tax rates are higher; in fact, the point estimates suggest that subjects if anything spend more when taxes are a higher fraction of the after-tax price.

5. Learning Over the Course of the Experiment

It is possible that in this lab setting, subjects do not fully understand the tax system they face at the start of the experiment. Despite the fact that subjects received clear feedback at the end of each round, this could cause us to find a salience effect simply because it took the subjects in the TE group longer to understand that the taxes would be added on to the price at checkout . To test for this possibility, we allow the subjects time to learn about the experimental design in the first rounds of the experiment. Column 1 of Table 11 drops the first half of the rounds of Part I while column 2 drops the first quarter of the rounds of Part I and Part II. Our results are similar to those presented in the main analysis; if anything, our results show a slightly larger salience effect, especially at the low tax rate.

C. Other Analyses

1. Time Spent per Round

The results from the previous section suggest that subjects are less attentive to low salience taxes. One reason subjects may respond more to taxes that are included in the displayed prices is that it is time-consuming or mentally costly to calculate the tax-inclusive price when taxes are excluded from the displayed price. Table 12 examines whether subjects spend more time deciding on what to purchase in rounds where taxes are not included in the displayed price. Column 1 pools subjects across both tax rates to determine the overall effect of tax-exclusive pricing on the duration of each purchasing decision. The average subject spends 12 seconds per round which does not differ

statistically when taxes are included versus excluded from the price. In fact, the point estimate on β_{TE} is negative suggesting that subjects spend slightly less time on their purchasing decisions when they are required to compute the after-tax price themselves. Column 2 tests for the effect of tax-exclusive pricing by tax rate. These results show that subjects spend the same amount of time on tax-inclusive and tax-exclusive rounds at both high and low tax rates.

These results may provide support for a heuristic model of decision-making – all subjects ignore taxes that are not included in the displayed price regardless of the size of the tax and, therefore, spend the same amount of time making their purchasing decisions when prices are tax-inclusive versus tax-exclusive. Therefore, if the duration of the purchasing decision is an accurate proxy for the cognitive cost required to pay attention to low salience taxes, these null results provide evidence against a bounded rationality model.

2. Number of Attempts

As mentioned in Section II, subjects were presented with a list of products on the first screen of the experiment then sent to a checkout screen that displayed the total after-tax price of their basket once they had selected which products they would like to purchase. However, after viewing their bill, they were able to return to the previous screen and adjust their purchases; in fact, customers whose basket exceeded the allotted budget of \$18 were forced to do so. If subjects facing tax-exclusive prices ignore the taxes when selecting their basket but realize they have over-spent once reaching the checkout page, we would expect to see an impact of tax-framing on the number of attempts a subject takes to decide on his final basket. Similarly, we might expect to see a larger salience effect if we limit our sample to subjects who did not alter their purchasing decision after reaching the checkout page.

Table 13 presents the results of this analysis. Column 1 presents results of the effect of the treatment on the number of attempts a subject makes per round. While some subjects took up to 10 attempts in a given round, the average subject took 1.3 attempts to make his final purchase. We find that facing tax-exclusive pricing had no effect on the number of attempts. Column 2 repeats our main analysis, but limits the sample only to rounds in which the subject did not adjust his purchases after reaching the checkout. Column 3 also repeats the main analysis, but uses a subject’s selected purchases from the first attempt, i.e., before making any adjustments after seeing

the checkout screen. We find that the salience effect at both tax rates increases in size in Columns 2 and 3, suggesting that allowing for adjustment after displaying the after-tax total amount may dampen the true salience effect.

3. Self-Reported Attention

Our salience parameter throughout this paper is measured by differences in expenditure when confronted with tax-exclusive versus tax-inclusive prices. However, this tax-framing may have an effect on whether an individual thinks about taxes when making a purchasing decision. In our post-experiment questionnaire, we ask subjects directly about whether they considered the taxes they faced when making their purchasing decisions. Table 14 compares this self-reported measure of attention for subjects in the treatment versus control groups. We find that subjects facing tax-exclusive prices in Part I (the treatment group) were significantly more likely to report considering taxes when making purchasing decision during those rounds. So while subjects in this group are more likely to overspend on their purchases, they are also more likely to be aware of taxes when making this decision.

V. Conclusion

Theories of bounded rationality predict that salience effects should fade as the utility costs of ignoring a tax grow larger for consumers. We find no evidence for this phenomenon over the tax rates we consider. Although we believe it likely that virtually all consumers would become attentive to a sufficiently large tax, our experimental results suggest the size of the tax that would trigger such behavior may exceed the magnitude of taxes that are currently levied in the United States and that are usually considered policy-relevant. In addition, to the extent we document a difference in salience between the high and low taxes we consider, our findings suggest consumers may be more attentive to the lower tax than to the higher one – a pattern of behavior that can be explained by a model of confirmation bias.

Our results are subject to several limitations. First, the usual caveats for interpreting laboratory experiments apply here. The biases present in our sample may differ from those in the general population (although one would probably not expect salience biases to be stronger in our sample

of undergraduates from a selective institution than in the overall population). Second, although we took a number of steps to make the purchasing decision in the experiment as realistic as possible (e.g., simulating an online shopping environment, providing real stakes and opportunities to correct mistakes), subjects might behave differently when making purchasing decisions in other contexts. Along these lines, although the economic stakes of the experiment are typical of the real-world purchasing decisions we attempt to emulate, the incentive for an individual to adopt a rule of always accounting for taxes is weaker here than under a real tax rate change because the experimental purchasing decision is not repeated once the experiment concludes. Finally, we wish to emphasize that our strongest result is the finding that salience effects persist even at the high tax rate we consider; the evidence we present in favor of confirmation bias is primarily suggestive. Additional research that can more precisely distinguish between competing positive models would be valuable.

Disclosures

The authors received funding for this project from the Princeton University Industrial Relations Section and from the Princeton Lab for Experimental Social Science (PLESS).

References

- Bradley, Sebastien.** 2017. “Inattention to Deferred Increases in Tax Bases: How Michigan Homebuyers are Paying for Assessment Limits.” *Review of Economics and Statistics*, 99(1): 53–66.
- Bradley, Sebastien, and Naomi Feldman.** 2015. “Hidden Baggage: Air Traveler Responses to Add-On Taxes and Fees.” *Working Paper*.
- Brown, Jennifer, Tanjim Hossain, and John Morgan.** 2010. “Shrouded Attributes and Information Suppression: Evidence from the Field.” *The Quarterly Journal of Economics*, 125(1): 859–876.
- Brown, Zachary, Nick Johnstone, Ivan Haščič, Laura Vong, and Francis Barascud.** 2013. “Testing the Effect of Defaults on the Thermostat Settings of OECD Employees.” *Energy Economics*, 39: 128–134.
- Camerer, Colin.** 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- Camerer, Collin, and Robin Hogarth.** 1999. “The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework.” *Journal of Risk and Uncertainty*, 19(1-3): 7–42.
- Chetty, Raj, Adam Looney, and Kory Kroft.** 2007. “Salience and Taxation: Theory and Evidence.” *National Bureau of Economic Research 13330*.
- Chetty, Raj, Adam Looney, and Kory Kroft.** 2009. “Salience and Taxation: Theory and Evidence.” *American Economic Review*, 99(4): 1145–1177.
- Chetty, Raj, John N Friedman, Tore Olsen, and Luigi Pistaferri.** 2011. “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records.” *The Quarterly Journal of Economics*, 126(2): 749.
- Colantuoni, Francesca, and Christian Rojas.** 2015. “The Impact of Soda Sales Taxes on Consumption: Evidence from Scanner Data.” *Contemporary Economic Policy*, 33(4): 714–734.

- Feldman, Naomi, and Bradley Ruffle.** 2015. “The Impact of Including, Adding and Subtracting a Tax on Demand.” *American Economic Journal: Economic Policy*, 2(1): 95–118.
- Finkelstein, Amy.** 2009. “E-Ztax: Tax Salience and Tax Rates.” *The Quarterly Journal of Economics*, 124(3): 969–1010.
- Gallagher, Kelly S., and Erich Muehlegger.** 2011. “Giving Green to Get Green? Incentives and Consumer Adoption of Hybrid Vehicle Technology.” *Journal of Environmental Economics and Management*, 61(1): 1–15.
- Goldin, Jacob.** 2015. “Optimal tax salience.” *Journal of Public Economics*, 131: 115–123.
- Goldin, Jacob, and Tatiana Homonoff.** 2013. “Smoke Gets in your Eyes: Cigarette Tax Salience and Regressivity.” *American Economic Journal: Economic Policy*, 5(1): 302–336.
- Haggag, Kareem, and Giovanni Paci.** 2014. “Default Tips.” *American Economic Journal: Applied Economics*, 6(3): 1–19.
- Hart, William, Dolores Albarracín, Alice Eagly, Inge Brechan, Matthew Lindberg, and Lisa Merrill.** 2009. “Feeling Validated Versus Being Correct: A Meta-Analysis of Selective Exposure to Information.” *Psychology Bulletin*, 4(135): 555–588.
- Hayashi, Andrew, Brent Nakamura, and David Gamage.** 2013. “Experimental Evidence of Tax Salience and the Labor-Leisure Decision Anchoring, Tax Aversion, or Complexity?” *Public Finance Review*, 41(2): 203–226.
- Hayashi, Andrew T.** 2014. “The Legal Salience of Taxation.” *The University of Chicago Law Review*, 1443–1507.
- Nickerson, Raymond S.** 1998. “Confirmation Bias: A Ubiquitous Phenomenon in Many Guises.” *Review of General Psychology*, 2(2): 175–220.
- Reck, Daniel.** 2015. “Taxes and Mistakes: What’s in a Sufficient Statistic?” *University of Michigan mimeo*.
- Taubinsky, Dmitry, and Alex Rees-Jones.** 2016. “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment.” National Bureau of Economic Research 22545.

Table 1: Displayed Product Prices by Tax Rate Group

	8 Percent Tax Rate		22 Percent Tax Rate	
	Tax-Exclusive	Tax-Inclusive	Tax-Exclusive	Tax-Inclusive
Nutella	3.23	3.49	2.86	3.49
Starbucks Via Instant Coffee (Box of 12)	3.95	4.27	3.50	4.27
Coconut Water	0.94	1.02	0.84	1.02
Trader Joe’s Trail Mix	2.92	3.15	2.58	3.15
Clif Bars (Box of 12)	4.87	5.26	4.31	5.26
Ferrero Rocher Chocolate Truffles	5.22	5.64	4.62	5.64
Pilot G-2 Gel Pens (3-Pack)	2.03	2.20	1.80	2.20
5-Hour Energy Drink (3-Pack)	3.40	3.67	3.01	3.67
Laundry Detergent (50 oz)	3.59	3.88	3.18	3.88
Pringles	1.06	1.14	0.93	1.14

Displayed prices are the prices subjects see prior to checkout.

Treatment group participants are shown tax-exclusive prices in Part I of the experiment and tax-inclusive prices in Part II.

Control group participants are shown tax-inclusive prices in both Part I and Part II.

Table 2: Pre-Tax Product Prices by Tax Rate Group

	8 Percent Tax Rate		22 Percent Tax Rate	
	Low Price	High Price	Low Price	High Price
Nutella	3.23	3.49	3.23	4.26
Starbucks Via Instant Coffee (Box of 12)	3.95	4.27	3.95	5.21
Coconut Water	0.94	1.02	0.94	1.24
Trader Joe’s Trail Mix	2.92	3.15	2.92	3.84
Clif Bars (Box of 12)	4.87	5.26	4.87	6.42
Ferrero Rocher Chocolate Truffles	5.22	5.64	5.22	6.88
Pilot G-2 Gel Pens (3-Pack)	2.03	2.20	2.03	2.68
5-Hour Energy Drink (3-Pack)	3.40	3.67	3.40	4.48
Laundry Detergent (50 oz)	3.59	3.88	3.59	4.73
Pringles	1.06	1.14	1.06	1.39

Low prices, which are consistent across tax rates, are used for the majority of analyses.

High prices are used in Section IV.B.2.

Table 3: Randomization Checks

	Age (1)	Male (%) (2)	White (%) (3)	High-Income (%) (4)	Upperclassmen (%) (5)	# Subjects (6)
Control - 8%	20.2 (1.2)	27.9 (45.4)	39.5 (49.5)	34.9 (48.2)	72.1 (45.4)	43
Control - 22%	20.5 (1.5)	32.4 (47.5)	40.5 (49.8)	45.9 (50.5)	75.7 (43.5)	37
Treatment - 8%	20.1 (1.2)	31.7 (46.8)	51.2 (50.3)	56.1 (49.9)	68.3 (46.8)	82
Treatment - 22%	20.6 (3.4)	29.2 (45.8)	52.3 (50.3)	53.8 (50.2)	67.7 (47.1)	65
Overall	20.3 (2.1)	30.4 (46.1)	47.6 (50.1)	49.8 (50.1)	70.0 (45.9)	227
Prob>F	0.55	0.96	0.42	0.12	0.82	

P-value is associated with the F-test to reject equality between treatment groups.

“High-income” indicates students who self-identified as upper-class or wealthy.

“Upperclassmen” indicates juniors or seniors in college.

Table 4: Descriptive Statistics

	Overall
Expenditures Per Round (\$)	
Mean	6.16
Standard Deviation	6.15
Median	5.26
Fraction of Rounds (%)	
with Expenditures > \$0	67.1
with Expenditures > \$15	16.6
Fraction of Individuals (%)	
with Expenditures > \$0 in Any Rounds	87.2
with Expenditures > \$0 in All Rounds	43.1
with Expenditures > \$15 in All Rounds	6.6
<i>N</i>	227

Table 5: Average Expenditure in Part I & II by Experimental Group

	Part I	Part II
	(1)	(2)
22 Percent Tax		
Treatment	7.23 (0.39)	5.95 (0.39)
Control	6.23 (0.51)	5.82 (0.48)
8 Percent Tax		
Treatment	6.42 (0.34)	6.03 (0.35)
Control	5.50 (0.43)	5.47 (0.44)
<i>N</i>	906	908

Table displays unadjusted means;
standard deviations in parentheses.

Table 6: The Effect of Tax-Inclusive and Tax-Exclusive Prices on Demand

	(1)	(2)
	Pooled	By Tax Rate
Tax-Exclusive	0.5694** (0.2222)	
Tax-Exclusive (8 Percent)		0.1741 (0.2453)
Tax-Exclusive (22 Percent)		1.0726*** (0.2880)
Round Fixed Effects	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes
F-stat		9.06
Prob>F		0.003
<i>N</i>	1,814	1,814

Standard errors clustered by individual reported in parentheses.

Outcome=total expenditures per round.

Analysis is a within-subject comparison of expenditure in Part I versus Part II.

F-test of equality between the tax-exclusive coefficient for 8 vs. 22 percent.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 7: Robustness: Downward-sloping Demand

	(1) Low Tax Group	(2) High Tax Group
8 Percent Price Increase	-0.5856*** (0.1318)	
22 Percent Price Increase		-0.9195*** (0.1740)
Round Fixed Effects	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes
<i>N</i>	1,343	1,111

Standard errors clustered by individual reported in parentheses.

Outcome=total expenditures per round.

Analysis is a within-subject comparison of expenditure in rounds with high versus low prices.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 8: Robustness: Censoring

	(1) Drop All Zero	(2) Drop All Max
Tax-Exclusive (8 Percent)	0.1927 (0.2742)	0.1917 (0.2568)
Tax-Exclusive (22 Percent)	1.2266*** (0.3191)	1.1439*** (0.3022)
Round Fixed Effects	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes
F-stat	9.71	9.26
prob>F	0.002	0.003
<i>N</i>	1,614	1,718

Standard errors clustered by individual reported in parentheses.

Outcome=total expenditures per round.

Analysis is a within-subject comparison of expenditure in Part I versus Part II.

Column 1: drops 25 subjects with zero purchases in every round.

Column 2: drops 12 subjects spending over 15 dollars in every round.

F-test of equality between the tax-exclusive coefficient for 8 vs. 22 percent.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 9: The Effect of Tax-Inclusive and Tax-Exclusive Prices on Demand, Extensive and Intensive Margin

	(1)	(2)
	Extensive	Intensive
Tax-Exclusive (8 Percent)	0.0474*	-0.0056
	(0.0273)	(0.0420)
Tax-Exclusive (22 Percent)	0.0811***	0.1650***
	(0.0294)	(0.0531)
Round Fixed Effects	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes
F-stat	1.23	8.67
prob>F	0.27	0.004
<i>N</i>	1,814	1,218

Standard errors clustered by individual reported in parentheses.

Column 1: outcome=indicator for positive expenditures per round.

Column 2: outcome=log expenditures per round for rounds with positive expenditures.

Analysis is a within-subject comparison of expenditure in Part I versus Part II.

F-test of equality between the tax-exclusive coefficient for 8 vs. 22 percent.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 10: Robustness: Tax Aversion

	(1)	(2)
	Controls Only	TI Rounds
22 Percent Tax Rate	1.4312	0.6269
	(1.1398)	(0.7065)
Round Fixed Effects	Yes	Yes
Menu Fixed Effects	Yes	Yes
<i>N</i>	639	1,227

Standard errors clustered by individual reported in parentheses.

Outcome=total expenditure per round.

Analysis is a between-subject comparison of expenditure by subjects.

randomly assigned to the high versus low tax rate groups.

Column 1: includes rounds with tax-inclusive prices for control group subjects only.

Column 2: includes all rounds with tax-inclusive prices.

Includes individual-level covariates: age, race, sex, income, and class year.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 11: Robustness: Learning

	(1)	(2)
Tax-Exclusive (8 Percent)	0.3866 (0.2609)	0.5397* (0.3092)
Tax-Exclusive (22 Percent)	1.0737*** (0.3385)	1.0938*** (0.3592)
Round Fixed Effects	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes
F-stat	4.30	2.71
prob>F	0.039	0.101
<i>N</i>	1,355	1,359

Standard errors clustered by individual reported in parentheses.

Outcome is total expenditures per round.

Analysis is a within-subject comparison of expenditure in Part I versus Part II.

Columns 1: drops the first 50% of rounds in Part I.

Columns 2: drops the first 25% of rounds in Part I and II.

F-test of equality between the tax-exclusive coefficient for 8 vs. 22 percent.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 12: Duration

	(1) Pooled	(2) By Tax Rate
Tax-Exclusive	-1.1276 (0.7886)	
Tax-Exclusive (8 Percent)		-0.8486 (0.9311)
Tax-Exclusive (22 Percent)		-1.4828 (1.0004)
Round Fixed Effects	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes
F-stat		0.33
prob>F		0.569
<i>N</i>	1,814	1,814

Standard errors clustered by individual reported in parentheses.

Outcome=time spent on purchasing decision per round in seconds.

Analysis is a within-subject comparison of duration in Part I versus Part II.

F-test of equality between the tax-exclusive coefficient for 8 vs. 22 percent.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 13: Attempts

	(1)	(2)	(3)
	Number of Attempts	Expenditure (One Attempt Only)	Expenditure (First Attempt)
Tax-Exclusive (8 Percent)	0.0030 (0.0687)	0.3584 (0.2573)	0.3952 (0.2812)
Tax-Exclusive (22 Percent)	0.0844 (0.0739)	1.4169*** (0.3229)	1.9003*** (0.3660)
Round Fixed Effects	Yes	Yes	Yes
Individual-by-Menu Fixed Effects	Yes	Yes	Yes
F-stat	1.02	10.52	15.44
prob>F	0.313	0.001	0.0001
<i>N</i>	1,814	1,550	1,814

Standard errors clustered by individual reported in parentheses.

Columns 1: outcome=number of attempts per round.

Columns 2: outcome=total expenditure per round for rounds with only one attempt.

Columns 3: outcome=total expenditure per round in first attempt.

Analysis is a within-subject comparison of the number of attempts or expenditure in Part I versus Part II.

F-test of equality between the tax-exclusive coefficient for 8 vs. 22 percent.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 14: Self-Reported Attention to Taxes

	Self-Reported Attention
Treatment (8 Percent)	0.3168*** (0.0829)
Treatment (22 Percent)	0.4315*** (0.0909)
High Tax Rate (22 Percent)	0.0040 (0.0974)
<i>N</i>	227

Standard errors clustered by individual reported in parentheses.

Outcome=indicator for self-reporting having considered taxes in Part I.

Analysis is a between-subject comparison of expenditure by subjects randomly assigned to each experimental group.

Includes individual-level covariates: age, race, sex, income, and class year.

{* $p < .10$, ** $p < .05$, *** $p < .01$ }

Figure 1: Instructions

(a) Treatment

There are 16 rounds in the experiment. In each round you will receive \$18, which you can spend on various products sold by the online store. The products are very cheap – the store is selling them at heavily-reduced prices. The prices at the store do not include tax - when you proceed to the checkout, an 8% tax will be added to your bill.

You can buy as many of each product as you like, but you cannot spend more than \$18 in one round. You do not have to spend your entire budget. Remember, you get \$18 to spend in each round.

At the end of the experiment, one of the rounds will be randomly selected. You will receive the products that you purchased in that round plus any money from that round that you did not spend.

Please make all decisions on your own without talking to anyone else. Based on your decisions, you can leave here with \$18 of cash and no products, only products and no cash, or some combination of cash and products.

(b) Control

There are 16 rounds in the experiment. In each round you will receive \$18, which you can spend on various products sold by the online store. The products are very cheap – the store is selling them at heavily-reduced prices. The prices at the store include an 8% tax – there is nothing extra added to your bill when you checkout.

You can buy as many of each product as you like, but you cannot spend more than \$18 in one round. You do not have to spend your entire budget. Remember, you get \$18 to spend in each round.

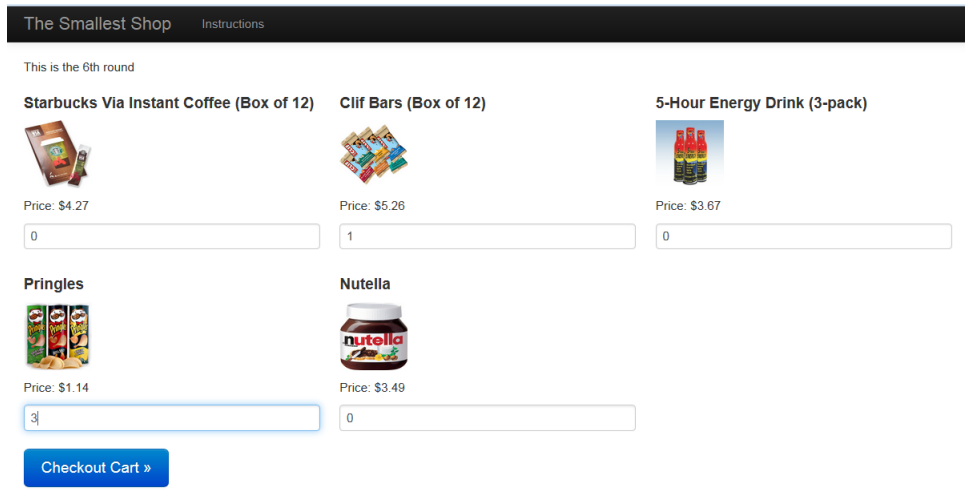
At the end of the experiment, one of the rounds will be randomly selected. You will receive the products that you purchased in that round plus any money from that round that you did not spend.

Please make all decisions on your own without talking to anyone else. Based on your decisions, you can leave here with \$18 of cash and no products, only products and no cash, or some combination of cash and products.

Note: Instructions are for subjects assigned to the low-tax groups.

For the high-tax groups, “8%” is replaced by “22%”.

Figure 2: Shopping Screen



Note: The products on this shopping screen are of one of the four possible product menus.

Sample prices above are tax-inclusive.

Figure 3: Checkout Screen




Figure 4: New Store Screen

A New Store

For the remaining rounds, you will be shopping at a different online store. The prices at this new store include an 8% tax – there is nothing extra added to your bill when you checkout.

As a reminder, at the end of the experiment one of the rounds will be randomly selected. You will receive the products that you purchased in that round plus any money from that round that you did not spend. The randomly-selected round may be from either the new store or the old store.

[Continue to Final Rounds »](#)

Note: Instructions are for subjects assigned to the low-tax groups.

For the high-tax groups, “8%” is replaced by “22%”.

Figure 5: Distribution of Expenditures Per Round

