PRICE ISN'T EVERYTHING: BEHAVIORAL RESPONSE AROUND CHANGES IN SIN TAXES

Alex Rees-Jones and Kyle Rozema

In traditional economic models, taxes change behavior by changing prices. In empirical analyses, factors other than price are thought to be relevant, but any nonprice factors are usually assumed to be held constant as taxes vary. We contend that violations of this assumption are expected when laws are passed changing sin taxes. In support of this claim, we document that state-level cigarette tax increases are concomitant with increases in antismoking appropriations, media coverage on smoking, lobbying efforts, and place-based smoking restrictions. The influence of these nonprice factors is easily confused with price effects, and we find evidence suggesting that controlling for them substantially reduces the estimated demand responsivity to the tax itself.

Keywords: sin taxes, legal changes, prices, expressive effects

JEL Codes: D9, H2, K4

I. INTRODUCTION

Taxes are often deployed to shape the behavior of individuals and institutions. But how does this behavioral change arise? Nearly universally in economic models, taxes' influence on behavior is attributed to their impact on prices. Concretely, a tax on a good raises its price, and the law of demand then implies that the quantity demanded should decrease. In the course of changing a tax law, however, more than just prices may change. Especially in cases where taxes are intended to dissuade consumption (e.g., taxes on alcohol, carbon, or cigarettes), changing the law typically requires making the case that consumption should be dissuaded. Interested parties may attempt to influence the legislative process with the provision of information

Alex Rees-Jones: The Wharton School, University of Pennsylvania and National Bureau of Economic Research, Philadelphia, PA, USA (alre@wharton.upenn.edu); Kyle Rozema: Washington University School of Law, St. Louis, MO, USA (kylerozema@wustl.edu)

Electronically published February 24, 2023

National Tax Journal, volume 76, number 1, March 2023.
 https://doi.org/10.1086/723093

and persuasive appeals, and they may attempt to simultaneously deploy other interventions to achieve their goal of dissuasion. Furthermore, the act of changing the law may itself have an *expressive effect*, whereby it directly influences beliefs, emotions, or behavior purely by condoning or condemning an activity (McAdams, 2015). In short, while tax changes do typically influence price, they may also be expected to occur alongside significant changes to potentially important nonprice factors.

In this article, we document the significant changes to nonprice factors that occur as cigarette taxes are changed. We further document that accounting for nonprice factors is critical in estimating the behavioral response to these tax changes. In our context, failure to control for our set of nonprice factors would lead a researcher to significantly inflate estimates related to tax elasticity. This omission of controls could also lead a researcher to infer that behavioral response preceding a tax change is unambiguous evidence of the anticipatory response present in rational-addiction models, when it may instead be driven by the coevolution of nonprice factors.

In our empirical analysis, we specifically examine how the consumption of cigarettes by pregnant women evolves throughout the process of state-level cigarette tax law changes. Dissuading smoking by pregnant women is of particular policy interest because of smoking's consequences for infant health (Evans, Ringel, and Stech, 1999; DeCicca, Kenkel, and Lovenheim, 2022). Prominent papers on the response to sin taxation use this behavior as a proxy for smoking more broadly (see, e.g., Gruber and Kőszegi, 2001), in part because of the unusual availability of high-frequency measurements of smoking that are comparable across states and over time. These data come from the smoking information that is reported on the US Standard Certificate of Live Birth, a permanent legal record required for all live births. Using these data, we examine how cigarette consumption evolves in the time period surrounding 150 state-level tax law changes occurring between 1989 and 2009.

Our empirical analysis proceeds in three steps. First, we examine the time paths of four nonprice factors: tobacco industry spending on political donations to state politicians, antismoking appropriations (which are used to fund antismoking advertising, among other things), news headlines related to cigarettes, and place-based legal restrictions on cigarette smoking (such as bans on smoking in restaurants or workplaces). We document substantial changes in these factors in the window surrounding a tax change. Compared with states not facing tax changes, states facing a change face a sizable increase in political donations and newspaper coverage in the lead-up to the reform. This illustrates that cigarette tax changes occur during time periods of substantial information provision and attempts at persuasion. The period before and after a cigarette tax change is also marked by an increase in both antismoking appropriations and place-based legal restrictions, illustrating the simultaneous deployment of other dissuasive policies. These findings establish that potentially important nonprice factors are not held constant throughout the lawmaking process.

Second, we examine the role of nonprice factors as explanations of the decrease in cigarette consumption that occurs around changes in tax law. We demonstrate that this decrease in consumption is partly explained by increases in antismoking appropriations, increases in newspaper coverage, and the adoption of place-based legal restrictions, even after accounting for the changes in the taxes themselves. The inclusion of these controls leads to quantitatively important changes in the behavioral response attributed to the tax itself: it reduces the estimated responsivity to the tax by roughly half. Under reasonable assumptions of the importance of unobserved nonprice factors, the estimated responsivity decreases further.

This analysis demonstrates that the within-state variation in demand occurring around a tax change is significantly influenced by variation in nonprice factors even when accounting for variation in taxes themselves. Combined with the findings that nonprice factors change markedly in the window around a tax change, this finding poses a challenge for identification strategies that make use of within-state variation in smoking and taxes. Within-state variation is a common source of variation used for identification in the sin-tax literature, and we are unaware of any observational study of a sin-tax elasticity where nonprice factors are reasonably controlled. Caution is therefore needed in mapping estimates from the existing literature into notions of elasticity derived purely from price effects.

In our third set of analyses, we investigate additional predictions that arise if nonprice factors are at work. Because the increase in nonprice factors often occurs well in advance of the tax change, the decrease in demand due to these factors should occur well before taxes become relevant for prices. Directly comparing the time path of demand in states facing a tax change to that in states facing no tax change, we find that demand is better explained as a smooth and gradual decrease preceding the change than it is by a function with discontinuity when tax changes are announced or enacted. These patterns contrast with the predictions of standard models where consumers only respond to the current tax and with the predictions of rational addition models where consumers respond to announced future taxes. We also document that a classic piece of evidence in support of rational-addiction models — the anticipatory response to tax increases between the moment of announcement and enactment found by Gruber and Kőszegi (2001) — does not survive the inclusion of controls for our nonprice factors. Our analysis does not establish the absence of the anticipatory rational-addiction response discussed by Gruber and Kőszegi (2001) but rather makes clear that compellingly testing for its presence is difficult when accounting for nonprice factors.

In sum, our results support the view that the process of changing tax laws changes more than just prices. Perhaps surprisingly, the quantitative impact of informational, psychological, and sociological factors influencing demand may be of at least comparable quantitative importance for understanding behavior during a tax change. Some of the role of these nonprice factors falls under the umbrella of "expressive powers of the law" — that is, the power of the law or lawmaking process to influence behavior through what it expresses and communicates, and not merely through its direct influence on rules and incentives. Such expressive effects have a prominent role in the legal literature, but this concept is rarely considered in economic applications,[1] and relatively little empirical work measures its importance.[2] While few would disagree with the conceptual point that these factors need to be held constant in assessments of behavioral response, the results of this paper suggest that holding these factors constant is more challenging and more important than is commonly appreciated.[3]

The paper proceeds as follows. Section II provides a conceptual framework for understanding the role of nonprice factors. Section III describes the construction of our data set. Section IV presents empirical analysis. Section V concludes.

II. CONCEPTUAL FRAMEWORK

Our empirical analyses evaluate the extent to which changes in cigarette consumption around tax changes are attributed to changes in prices. We formalize the potential theoretical role of nonprice factors using the model of Reif (2019), which combines and generalizes the forward-looking features of the rational-addiction model of Becker and Murphy (1988) and social interaction features in the spirit of Brock and Durlauf (2001).

Consider an individual facing the following utility maximization problem:

$$\max_{\{a_t, c_t\}\forall t} \sum_{t=1}^{\infty} \beta^{t-1}(U(a_t, c_t, x_t, S_t) + G(a_t, E_t[\bar{a}_t]))$$

$$\text{s.t. } A_0 = \sum_{t=1}^{\infty}(1 + r)^{-(t-1)}(c_t + p_t a_t)$$

$$S_{t+1} = (1 - d)(S_t + a_t).$$

Utility in period t is governed by the individual's discount rate ($\beta < 1$), private utility (U), and social utility (G). Utility is maximized subject to a budget constraint requiring that lifetime discounted expenditures (at an interest rate r) equal lifetime discounted wealth (A_0).

[1] This is true not only in classical economic models with rational agents (e.g., Pigou, 1920; Ramsey, 1927; Harberger, 1964) but also in economic models of sin taxation in behavioral public finance (e.g., O'Donoghue and Rabin, 2003, 2005, 2006; Farhi and Gabaix, 2017; Allcott, Lockwood, and Taubinsky, 2019). Although these behavioral-economic models involve agents with some degree of "irrationality," they continue to model taxes as affecting quantity demanded through prices. A partial exception to this characterization occurs in models of "tax aversion" in which tax changes dissuade behavior by more than an equivalent price change due to dislike of the tax itself (Kessler and Norton, 2016).

[2] For some discussion and measurement of expressive effects, see, e.g., McAdams and Nadler (2005), Funk (2007), Wittlin (2011), Dwenger et al. (2016), and Fabbri and Hoeppner (2018).

[3] Our findings are reminiscent of DeCicca et al. (2008), who document a positive cross-sectional correlation between cigarette prices and antismoking sentiments and show that the correlation leads to the overestimation of price elasticities inferred from comparisons across states with different tax regimes. Our point is related but distinct: that a variety of factors relating to and shaping antismoking sentiments rapidly change within a state in the lead-up to a tax change. These findings are clearly related, but our results more directly inform empirical strategies based on within-state variation.

This model has four basic features. First, the individual chooses between smoking a_t at price p_t and other consumption c_t (with price normalized to 1), resulting in substitution between these goods governed by their relative prices. Standard sensitivity to contemporaneous prices arises. Second, the individual is influenced by addiction to nicotine (achieved by smoking). Addiction is modeled with a "stock" of addictive capital S_t that decays each period by the rate of depreciation $d \in (0, 1)$. The larger the addictive stock, the higher the marginal utility of smoking. Through this channel, the amount of nicotine consumption today will influence how much an individual wants to smoke in the future; therefore, prices that will be faced in the future become relevant to consumption decisions today. Third, the utility of smoking is influenced by the behavior of other individuals through social utility (G), such as where conformity to group smoking norms is valued or where smoking generates spillovers.[4] Social utility is governed by the relationship of the individual's smoking (a_t) and the individual's expectation of the average smoking of others ($E_t[\bar{a}_t]$).[5] Fourth, utility is influenced by a catch-all component for "nonprice factors" x_t, which accounts for the level of education, advertising, and other factors that influence both utility directly and utility achieved from smoking.[6]

In theoretical or empirical approaches using models of this variety, the term x_t is normally viewed as a confounding factor. For convenience, it is typically assumed to be held constant in the course of a tax change. We argue that a variety of factors that evolve in the course of a tax change is naturally accommodated by this term. We follow Reif (2019) and represent individual utility (after substituting optimal c_t into the equation) as:

$$\begin{aligned} V(a_t, x_t, S_t, E_t[\bar{a}_t]) = {} & -\frac{1}{2}\left(b_{aa}a_t^2 + b_{SS}S_t^2 + b_{xx}x_t^2 \right) + b_{aS}a_tS_t + b_{ax}a_tx_t + b_{Sx}S_tx_t \\ & + b_aa_t + b_SS_t + b_xx_t + b_k + G(a_t, E_t[\bar{a}_t]). \end{aligned}$$

This representation allows utility to depend on both linear and quadratic terms for smoking (a_t), addictive stock (S_t), and other factors (x_t). Utility also depends on interactions between all three, and a more general functional form of social utility ($G(a_t, E_t[\bar{a}_t])$). For further details and restrictions on these terms, see Reif (2019).

This representation highlights the manner in which nonprice factors may become relevant. In the context of this model, such factors may operate through several channels. First, through the terms b_{xx} and b_x, x_t directly affects utility. This could

4 As noted by Reif, social utility of these types have been employed in significant prior work (Binder and Pesaran, 2001; Brock and Durlauf, 2001; Glaeser and Scheinkman, 2002; Blanchflower, Van Landeghem, and Oswald, 2009).

5 Reif considers several ways of specifying this function and assumes that the group is sufficiently large that the individual's contribution to the mean is negligible.

6 For ease of exposition, we adopt Reif's treatment of x_t as a scalar. However, we note that its replacement with a vector poses no conceptual problem.

capture phenomena such as a direct aversion to graphic warning labels,[7] and it could capture direct aversion to antismoking sentiment (as in, e.g., DeCicca et al., 2008). Second, through the term b_{ax}, x_t influences the marginal utility generated by additional smoking. This could capture news coverage influencing beliefs about the marginal health consequences of smoking or direct increases in the marginal costs of smoking from non-tax dissuasive policies such as place-based restrictions. Third, through the term b_{Sx}, x_t influences the marginal utility of smoking generated by addictive stock. This could capture advertisements that remind the individual of smoking more often, policies that restrict exposure to smoking in public areas, or increased information about or accessibility of smoking cessation aids.[8]

The degree to which smokers actively choose to manage their addiction in a forward-looking manner remains a topic of debate. However, as documented in Reif (2019), convenient cigarette demand equations result from this framework regardless of whether consumers are myopic or forward-looking. The smoking demand equation for a myopic individual is:

$$a_t = \alpha^1 p_t + \alpha^2 S_t + \alpha^3 \bar{a}_t + \alpha^4 x_i + k_m.$$

The smoking demand equation for a forward-looking individual is:

$$a_t = \beta^1 p_t + \beta^2 p_{t+1} + \beta^3 S_t + \beta^4 \bar{a}_t + \beta^5 \bar{a}_{t+1} + \beta^6 x_t + \beta^7 x_{t+1} + k.$$

In both cases, demand is linear in the parameters of interest, with the coefficients (α, β) and the constants (k_m, k) depending on the specifics of the social interaction model.

These demand equations nest both the common considerations of a tax change present in the literature and the alternative nonprice factors that we set out to study. First, both demand equations contain a standard, contemporaneous price effect (represented by the term α^1 or β^1). In traditional models of sin taxation that are meant to approximate (addiction-free) rational agents, this is the sole channel through which a tax change influences behavior. Next, the demand equations contain intertemporal dependency on past and future prices. Past and future prices become indirectly relevant through the addictive component, captured by term α^2 in the myopic model or the term β^3 in the forward-looking model, and become directly

7 For recent empirical evidence on the effects of graphic warning labels on consumption, see Beleche et al. (2018). Graphic warning labels are used in other settings as well, and recent field-tests provide empirical evidence that graphic warning labels on sugary drinks can meaningfully decrease consumption (e.g., Roberto et al., 2016).

8 Reif's model does not permit these other factors to influence social utility. However, in principle, the importance of these factors could arise through this channel as well. Antismoking campaigns have arguably increased the social stigma surrounding smoking, which itself can decrease demand (Stuber, Galea, and Link, 2009; Riley et al., 2017). Furthermore, place-based smoking restrictions can displace smoking to or from public locations, so these restrictions could shape perceptions of smoking prevalence (Hamilton, Biener, and Brennan, 2008).

relevant for the forward-looking consumer through the term β^2. These terms capture the forces of anticipatory price responses (as in Becker and Murphy, 1988; Gruber and Kőszegi, 2001).

We contend that typical discussion or analysis of tax changes imagines taxes operating through only contemporaneous and non-contemporaneous prices. Indeed, most discussion and analysis of tax changes typically focuses entirely on contemporaneous prices. However, we draw attention to the remaining components of the demand equations, which are governed directly by issues such as expectations of social behavior, information, and material changes to the costs or benefits of smoking itself. A tax change is almost always viewed purely as a price change, but we will present evidence suggesting that both the social expectations and information components captured in the remaining terms of the equation are directly affected.

III. DATA SOURCES AND SAMPLE DEFINITION

Our data analysis relies on the cigarette consumption of pregnant mothers, the timing of state-level proposed tax law changes, and measurement of our candidate nonprice factors that could influence demand. We detail our construction of these data in this section.

A. Cigarette Demand: Natality Files

To conduct our analyses, we require a measure of cigarette consumption that is both high-frequency and comparable across states and over time. When facing these requirements, research on cigarette demand typically uses one of two data sources: direct measurement of cigarette sales or survey measures of consumption.

While sales data have clear appeal as a data source, they have three important limitations for the purposes of our study. First, rich sales data have limited availability for longer time series (e.g., Nielsen scanner data is not available for purchases before 2004). Second, cigarette purchasing diverges from cigarette consumption surrounding tax changes because of stockpiling (Chiou and Muehlegger, 2014). Third, tax changes may induce smokers to travel to other tax territories for purchases, which lead sales data from a particular retailer to more imperfectly reflect a given consumer's consumption (see, e.g., Chiou and Muehlegger, 2008; Lovenheim, 2008; DeCicca, Kenkel, and Liu, 2013a, 2013b).

Although survey measures of consumption can overcome these issues, they come with their own limitations. First, survey data are often criticized for selection into participation or imperfections in recorded responses. Second, few surveys are conducted in a manner that provides data on cigarette consumption at a granular time level.

To address these issues, we adopt the (partial) solution proposed by Gruber and Kőszegi (2001): we study the cigarette consumption of mothers who gave birth as recorded in the Vital Statistics Detailed Natality Data Files. The manner in which this data set is constructed mitigates concerns present with sales data. Certificates of Live Birth[9] have significantly fewer problems with noncompletion than typical surveys, and the formal records are made with input and supervision from medical professionals. The data are collected by first starting with medical records and physician reports of the mothers' smoking behavior. The questionnaire submitted by the mother is only completely relied upon in cases where further medical records are not available. When other medical records are available, a nurse follows up with the mother if they conflict with the questionnaire responses. Smulian et al. (2001) find that 86 percent of hospital staff in New Jersey maternity facilities used either prenatal care records or maternal hospital medical records as the source of the smoking information. They also found that only 6 percent of hospital staff in New Jersey used the mother's report.[10] Furthermore, the smoking measures included in these data elicit mothers' smoking behavior in the window prior to their child's birth. The fact that births are distributed across time, combined with the fact that these data were collected in a consistent way across a large group of states for a significant window of time, allows for the construction of a panel data set spanning many tax change events while still having the needed granularity.

For most states, the Natality files record every birth since 1989 and contain a standardized elicitation of recent cigarette consumption as of that year. We follow the treatment of the smoking variable in Gruber and Kőszegi (2001) and assume that it represents the average rate of consumption in the month before delivery. Unfortunately for our purposes, the details of the elicitation of cigarette consumption changed in the mid-2000s, and states transitioned to using a new measure in a staggered manner between 2003 and 2009. Because the new measures are not directly comparable to the old measures, we drop states from our data set once the older measure is no longer available. Moreover, information on cigarette consumption is either unavailable or unreliable in California, Indiana, South Dakota, and New York, so we follow Gruber and Kőszegi (2001) and drop these states. For the remaining states, we construct the state-year-month average cigarette consumption beginning in 1989 and ending when the survey elicitation changed in the state.

Like Gruber and Kőszegi (2001), we note that although pregnant women are not a representative population, they are an important group to study because of the important consequences of maternal smoking on infant health. Indeed, the impact of maternal smoking on infant health is thought to be one of the main externalities associated with smoking (Evans, Ringel, and Stech, 1999; DeCicca, Kenkel, and Lovenheim, 2022).

[9] See Figure A1 for an example of these forms.

[10] Of course, smoking behavior recorded in earlier medical records is also subject to concerns of misreporting — for example, if the mother misrepresented her smoking behavior in prenatal care visits. However, the fact that answers are checked for consistency over time, and the fact that the process is supervised by healthcare professionals and based on information from longer clinical relationships with the mothers, provides greater assurances of accuracy than simple, one-shot survey elicitations.

B. Timing of Tax Laws

For data on monthly state cigarette taxes in place and the timing of changes in state cigarette tax law, we use data from the State Tobacco Activities Tracking and Evaluation System from the Centers for Disease Control and Prevention (2018). We identify the timing of each tax law change, including year-month of the enactment of the new law and the year-month that the law change becomes effective. We also validate the effective dates by comparing them to those reported in the Tax Burden on Tobacco dataset released by the Federation of Tax Administrators.

C. Intensity of Political Activities and Social Debate

We form three different measures of political and social debate to serve as some of our nonprice factors. We view these as measures of activities intended to inform or influence the behaviors of either voters or lawmakers. These measures can be thought of as imperfect proxies for an expanded definition of the “expressive effects” (McAdams, 2015) — effects that occur as a result of the ideas communicated by the activities and debate that surround and accompany the process of legal change.

First, we establish a proxy relevant to the intensity of political debate around cigarettes using tobacco industry spending on political donations to state politicians. To do so, we use a component of the Database on Ideology, Money in Politics, and Elections (DIME) (Bonica, 2016). DIME contains information on approximately 100 million donations made by individuals, political action committees, and corporations to candidates in local, state, and federal elections from 1979 to 2014.[11] Our measure draws from the “contributions database” within DIME, focusing specifically on donations to state-level candidates. DIME includes the date on which the donation was made and the state of the recipient candidate or committee. Using these data, we build a panel of state-year-month donations made from the tobacco-related entities. It is worth noting that political donations are likely less salient to smokers than the other nonprice factors, in part because they are an input into a policy outcome rather than a policy outcome itself.

Second, to construct a measure of persuasive activity by antitobacco interests, we use state-level antismoking appropriations from Centers for Disease Control’s Health Communication Interventions (2015).[12] The appropriations data capture funding from four major funding sources: federal funding, state funding, the Robert Wood Johnson Foundation, and the American Legacy Foundation. Note that this measure of appropriations only imperfectly reflects expenditures in a given time period because the appropriations are not necessarily expended. We assume that they are expended in a uniform manner throughout the funding period.

[11] Within this database, contribution records, candidate and committee filings, and election outcomes for state elections are provided by the National Institute on Money in State Politics and the Sunlight Foundation.

[12] Because the appropriations data are available only since 1991, we code appropriations in 1989 and 1990 as those in 1991.

Third, we establish a proxy for the intensity of public debate around cigarettes by using headline appearances of smoking-related words in newspapers. This approach derives a measure of public debate by assuming that newspapers either respond to public demand for issues or promote debate about an issue by writing about it.[13] To construct the data set, we scraped headlines from newslibrary.com from 1990 to 2015. According to newslibrary.com, it is “the most complete archive available” for 6,504 newspapers and other news sources. The website is searchable by state and search results contain the date of the news article. Our main measure is the number of monthly headlines including “cigarette.” Newspapers enter and exit this data set over time, with substantially more newspapers in the data set in recent years, so we measure the intensity of debate using the number of cigarette-related headlines per state or local journal rather than the raw count of articles appearing.

D. Place-Based Legal Restrictions

We construct an index of the legislative activity related to place-based restrictions on smoking activity using the State Tobacco Activities Tracking and Evaluation System from the Centers for Disease Control and Prevention (2018). These data contain historical information on all state-level legislation, including the enactment date and the effective date. More important, for a set of classes of locations (including restaurants, private worksites, day care centers, and more),[14] this data set provides the citation of the statute that prohibited smoking at that location class. Our place-based legal restriction index is the average of the cumulative number of legal changes in each of the classes of locations that are independently tracked.

E. Summary of Panel

The final sample consists of 10,022 state-year-month cells and spans 150 tax changes occurring in 43 of the 46 states in the sample. Figure 1 reports the window

[13] In showing that changes in physician behavior before tort reforms lead to a twofold increase on estimated effects of the law on physician behavior, Malani and Reif (2015) document an increase in newspapers discussing medical malpractice reforms before it was adopted. Past work like this partially motivates our approach.

[14] The full set of classes of locations recorded in the data is “Bars, Commercial Day Care Centers, Government Multi-Unit Housing, Government Worksites, Home-Based Day Care Centers, Hotels and Motels, Personal Vehicles, Private Multi-Unit Housing, Private Worksites, Restaurants, Bingo Halls, Casinos, Enclosed Arenas, Grocery Stores, Hospitals, Hospital Campuses, Malls, Mental Health Outpatient and Residential Facilities, Prisons, Public Transportation, Racetrack Casinos, Substance Abuse Outpatient and Residential Facilities.” When counting the number of distinct changes, we group several similar classes. Our coding of day-care centers includes commercial and home-based facilities. Our coding of casinos includes casinos, bingo halls, and racetracks. Our coding of hospitals includes hospitals and hospital campuses. Our coding of mental health facilities includes mental health outpatient and residential facilities. Our coding of substance use facilities includes outpatient and residential facilities.

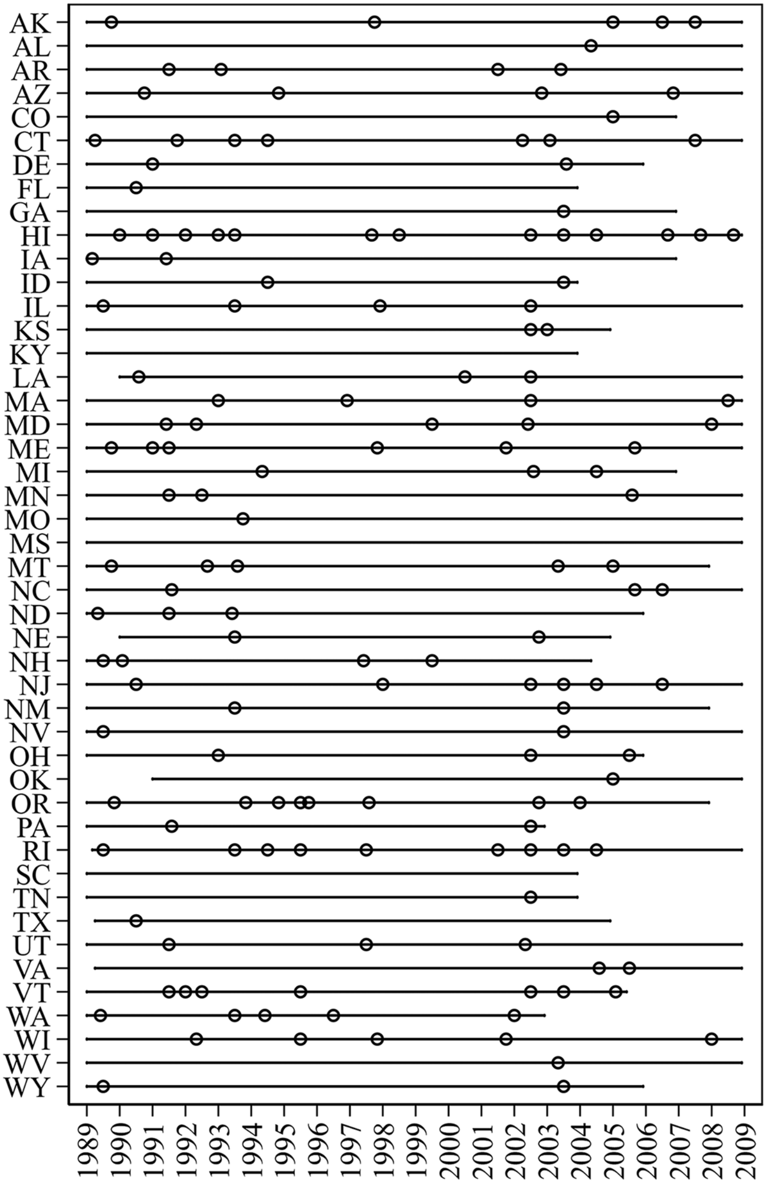

Figure 1. Summary of state's data availability and tax law changes. Each line indicates the years that the state is in the sample. The circle markers indicate the timing of cigarette tax changes. The dates for changes in state cigarette tax laws come from the State Tobacco Activities Tracking and Evaluation System from the Centers for Disease Control and Prevention (2018).

of available data for each state, as well as the timing of tax law change events. The data start in 1990 for all but a few states. The data end in 2009 for 23 states, the data end between 2005 and 2009 for 12 states, and the data end from 2003 to 2004 for the remaining 11 states. The size of the tax change ranges from a decrease of \$0.10 to an increase of \$1.00, with 75 percent of tax changes increasing by at least \$0.05, 50 percent increasing by at least \$0.15, and 25 percent increasing by at least \$0.37. The average size of the tax change was an increase of \$0.25.

IV. EMPIRICAL ANALYSIS

In this section, we investigate the importance of nonprice factors in explaining the decrease in cigarette consumption around tax law changes. We proceed in three steps. First, we examine the time path of nonprice factors around tax law changes. Second, we examine the relative importance of price and nonprice factors in explaining the decrease in consumption that occurs following tax law changes. Third, we revisit typical explanations for the timing of behavioral changes in the period surrounding a tax law change. The formal econometric specification varies across these three groups of analyses, and we present the relevant details when each specification is introduced. However, all three groups of analyses involve a "stacked-event-study" approach, which we describe here. For recent papers organizing data with similar procedures, see Cengiz et al. (2019) or Deshpande and Li (2019).

To illustrate the organization of data in a stacked event study, we first describe the data structure of a single event considered in isolation. We then discuss the process whereby individual events are "stacked."

Imagine that we are interested in understanding the impact of a single event: the tax change that occurred in Louisiana in July 2000. To measure the impact of this event, we would examine the evolution of smoking behavior in Louisiana in a window of time around that change, which we call an *event window*. Depending on our analysis, the event window may be divided into further subperiods. In analyses geared toward modeling the evolution of nonprice factors and consumption in the time before and after a tax change, we define a baseline period occurring well before the tax change and then examine how variables deviate from the baseline average in the time period more immediately before and after the tax change. Concretely, for our primary specifications in Subsections IV.A and IV.C, we define the window 25–36 months (two to three years) before the tax change as the baseline period and then examine how variables evolve relative to the average value during baseline period across the 24 months leading up to the tax change and the 12 months following it. In our regression analyses of tax responsivity, which involve comparing average cigarette consumption before and after the tax change, we will define the event window by a symmetric period before and after the tax change. Concretely, for our primary specification in Subsection IV.B, we consider an event window covering 36 months before the tax change and 36 months after the tax change (although we also consider narrower and wider windows).

Table 1

Descriptive Statistics

		Stacked Events		
	(1) Panel	(2) All	(3) Control	(4) Treatment
Average cigarettes per day	1.63	1.72	1.73	1.49
Monthly tobacco industry political donations per 1000 citizens	1.45	0.56	0.51	1.35
Monthly antismoking appropriations per 1000 citizens	131.65	100.90	94.42	190.19
Monthly news headlines per journal	0.36	0.35	0.35	0.35
Place-based legal restrictions index	0.44	0.34	0.33	0.47
Observations	10,022	78,204	72,912	5,292
Events	150	108	108	108

Note: Column 1 presents averages of our primary variables of interest in our panel data. Columns 2–4 present averages from our "stacked-event-study" data set applying an event window covering 36 months before the tax change to 12 months after.

Conceptually, the behavior observed in Louisiana during the event window may be thought of as a demonstration of behavior in a state that has been "treated" with a tax change. In some of our analyses, we focus on understanding the predictors of within-treated-state variation in outcomes. In other analyses, we compare changes in treated states with changes in "control" states that did not experience a tax change. In forming our set of control states, we exclude any state that experienced a tax change during our defined event window. In addition, nonprice factors begin to evolve for a period before a tax change and some continue to evolve after, so we exclude any control state that experienced a tax change in the six months immediately before or after the event window.[15]

For this single event, our event data set would consist of observations for Louisiana and for each control state for every month in the event window. To form our full "stacked-event-study" data set, we append the event data sets for all of the tax changes considered in our analysis. For our graphical analyses, and in Table 1, we will restrict the data set to the 108 events for which we observe all state-year-months in the full event window (effectively excluding tax changes at the very beginning or very end of the sampling periods plotted in Figure 1). In our regression analyses, we will consider all 150 tax change events.

Table 1 reports descriptive statistics of the original panel data and the stacked-event-study data set. The stacked-event-study data set analyzed here is constructed with an event window ranging from 36 months before the tax change to 12 months.

[15] While we believe this additional exclusion based on a six-month buffer is conceptually important, in practice it has little impact on the estimates that we present.

This data set contains 5,292 event-state-year-month cells for the treated group and 74,774 event-state-year-month cells for the control group. As this table reveals, treated states in the stacked-event-study data set differ in clear ways from both their control states and from the overall panel average. Relative to control states, treated states have lower average cigarette consumption and higher average tobacco industry political donations, antismoking appropriations, and place-based legal restrictions indices. These differences are consistent with the general concern motivating our paper: that tax changes do not occur at random but rather during periods when a variety of activities related to the dissuasion of smoking occur. In the next section, we begin our analysis by examining these differences more formally.

A. Examining the Time Paths of Nonprice Factors

We begin by examining the time path of our candidate nonprice factors around tax law changes. Within each event, we compare the evolution of each measure in the state facing a tax law change (the "treated" state) with states facing no tax law change within that window (or in the six months before and after it) (the "control" states). To assess this evolution, we calculate the average value of each measure in the baseline period (25–36 months before the tax change, indicated with shading in Figure 2) and then plot the deviation from this baseline average across the event window using a six-month moving average.

Using this methodology, Figure 2 documents a stark escalation of the four nonprice factors in the time surrounding a tax law change. Panel A reports the tobacco industry political donations, Panel B reports the antismoking appropriations, Panel C reports the news headlines per journal, and Panel D reports the place-based legal restrictions index. For a particularly clear example of the nature of social and political discourse in the course of a tax law change, consider the time path for tobacco industry political donations. While the baseline level of donations is quite low, donation activity spikes in a narrow window leading up to the tax law change. The other measures exhibit conceptually similar differences between treated and control states, although the changes are less localized. In the time period leading up to the tax law change, there is a marked increase in newspaper headlines concerning cigarettes and in antismoking appropriations. In the 24 months leading up to a tax law change, there is a steady divergence between the treated and control states in the place-based legal restrictions index.

Moving beyond visual assessment to statistical analysis, we directly test for the significance of these patterns in a difference-in-differences approach. We estimate the following regression with the stacked-event-study data:

$$x_{e,s,t} = \sum_{b \in B} \{\beta_b \times I(t \in b) + \gamma_b \times I(t \in b) \times T_{e,s}\} + \phi_{e,s} + \epsilon_{e,s,t} \tag{1}$$

for event e in state s and event-time t, where b denotes an event-time period bin and where B denotes the set of bins. In this equation, the outcome of interest is the value

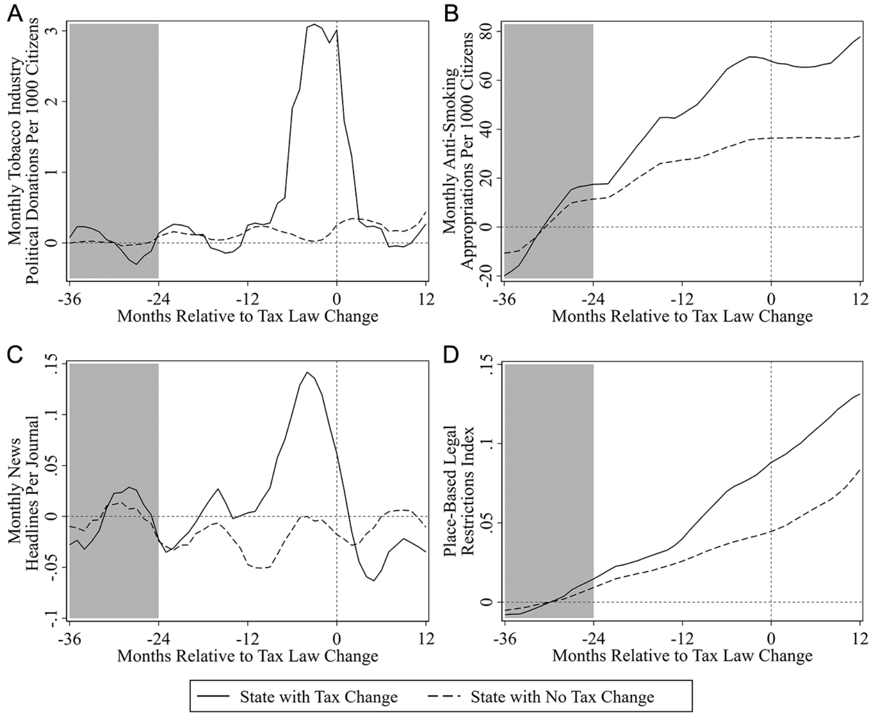

Figure 2. Time paths of nonprice factors. (A) Tobacco industry political donations. (B) Antismoking appropriations. (C) News headlines. (D) Place-based legal restrictions. Each panel reports the evolution of a nonprice factor, comparing "treatment" states with a tax change to "control" states with no tax change. We plot a six-month moving average, shifted such that the average of the nonprice factor in the shaded window 25–36 months before the tax change is zero. The data on tobacco industry spending on political donations to state politicians come from the Database on Ideology, Money in Politics, and Elections (DIME) (Bonica, 2016). The data on state-level antismoking appropriations come from Centers for Disease Control's Health Communication Interventions (2015). The data on appearances of smoking-related words in newspapers come from headlines in newslibrary.com. The data on place-based restrictions come from the State Tobacco Activities Tracking and Evaluation System from the Centers for Disease Control and Prevention (2018).

of a nonprice factor $x_{e,s,t}$ measured for event e in state s and event-time t. Variable $T_{e,s}$ takes the value of 1 if state s is the treated state for event e and is otherwise 0. Event time is grouped into six-month bins, spanning from 24 months prior to the event's tax change to 12 months after the events' tax change. The function $I(t \in b)$ takes the value of 1 if the event time under consideration falls in bin b and is otherwise 0. The term $\phi_{e,s}$ denotes an event-state fixed effect, measuring the average value of the nonprice factor during the event-and-state-specific baseline period (25–36 months prior to the event's tax change). For each six-month bin, this regression estimates the average difference in the value of the nonprice factor relative to its average value in the baseline period. For control states, this difference is captured by the term β_b. The term γ_b denotes the difference in these differences between treatment and control states. These serve as the key difference-in-difference estimates of interest.

Statistical inference in this framework requires accounting for several dimensions of correlation in the error terms. First, because of the manner in which events are stacked, a given state-year-month can be present multiple times in our data. For example, Alabama in January 2000 serves as a control observation for both Louisiana's July 2000 tax change and Arkansas's July 2001 tax change (as well as other tax changes).[16] This introduces correlation between the error terms for these two control observations in the data. Clustering standard errors by state accommodates this issue while additionally accounting for residual within-state variation occurring across events and across time.[17] Second, there can be common shocks experienced by all states at given moments in time, such as those from an increase in the federal cigarette tax or the death of the Marlboro Man, motivating us to additionally cluster standard errors by year-month. Finally, one might be concerned about correlation occurring across all observations within a given event, as can arise due to event-level selection of control states. This motivates us to additionally cluster at the event level. We perform multiway clustering using the approach of Correia (2016). We maintain this set of clustering practices for all analyses.

Figure 3 plots the estimated coefficients with 90 percent confidence intervals. The figure reveals that the patterns described above are generally statistically significant. Compared with control states, states facing a tax law change experience an increase in newspaper headlines in the lead-up to the date of the tax law change, an increase in place-based legal restrictions in the years following the tax law change, and a (statistically insignificant) increase in antismoking appropriations throughout. The increase in tobacco industry political donations in the immediate lead-up to the tax law change remains apparent, but the estimate does not reach statistical significance at conventional levels.

In summary, states experiencing a tax change also experience evolutions of related nonprice factors, both when examined in isolation and also when benchmarked against control states.[18] If these nonprice factors have their own direct influence on consumption, this demonstrates both that this evolution must be

[16] Note that the event time in which a given state-year-month will be a control observation differs across events. For example, Alabama in January 2000 serves as a control observation in event time -6 for Louisiana's July 2000 tax change and event time -18 for Arizona's July 2001 tax change.

[17] Note that to deal with the issue of repeated cases of the same observation, it would be sufficient to cluster at the state-year-month level. However, this level of clustering is nested within the state level, and clustering at this coarser level is standard practice in this literature.

[18] It is worth emphasizing that the apparent pretrends observed in Figures 2 and 3 are not spuriously generated by earlier tax changes included in the event data. Similar patterns arise when Figure 3 is estimated from the subset of events with no other tax changes occurring within the event window (see Figure A2).

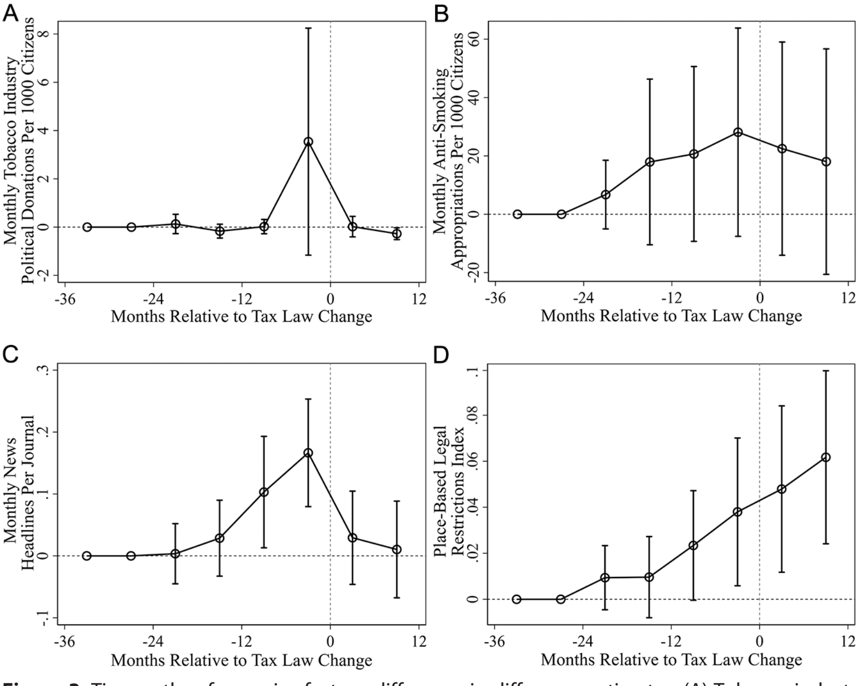

Figure 3. Time paths of nonprice factors: difference-in-difference estimates. (A) Tobacco industry political donations. (B) Antismoking appropriations. (C) News headlines. (D) Place-based legal restrictions. This figure reports the estimated coefficients associated with the difference-in-differences model presented in Equation (1). Capped lines indicate 90 percent confidence intervals.

controlled for when analyzing the behavioral response to tax changes and that it cannot be controlled for by simple comparisons to nontreated states.

B. Testing for Impact of Nonprice Factors

Having established that nonprice factors are not constant around tax law changes, we now assess whether this variation is relevant for understanding the change in cigarette consumption that happens within treated states as their laws change. To do so, we construct event windows with the same number of periods before and after a tax change and estimate how cigarette consumption over these windows responds to the change in cigarette taxes. We first consider a specification in which nonprice factors are not controlled, representing the baseline analysis that would be done when the researcher does not have these data. We then add our nonprice factors as controls and examine the consequences. Because we do not have exogenous variation in nonprice factors, and because we view our nonprice factor variables as an incomplete set of noisy proxies for the broader phenomenon of interest, we do not focus attention on their estimated coefficients and do not interpret their coefficients as causal estimates. Instead, we examine the impact of the inclusion of this partial set of controls on the estimated coefficients on the state tax.

To this end, in Table 2, we estimate versions of the following regression with stacked-event-study data:

$$\ln(\text{cigarettes})_{e,s,t} = \beta\,\text{statetax}_{e,s,t} + \gamma x_{e,s,t} + \phi_{e,s} + \epsilon_{e,s,t}. \tag{2}$$

Subscripts again denote event e in state s and event-time t. In words, we regress the natural log of cigarette consumption on the current state tax, nonprice factors $x_{e,s,t}$, and state-event fixed effects ($\phi_{s,e}$). Because we are interested in understanding the determinants of within-treated-state covariation in cigarette consumption and taxes, we restrict our data to only the treated state for each event.[19] This restrictions makes the state-event fixed effect equivalent to an event fixed effect. We initially estimate this model using an event window that contains 36 months of data both before and after the event. After presenting this initial model, we explore robustness to shorter and longer event windows.

Column 1 of Table 2 illustrates a well-documented finding: consistent with the existence of price responsivity, cigarette consumption is negatively associated with the state tax within the state facing a tax law change. Interpreting the magnitude of coefficients, a one dollar increase in the state tax is associated with a 0.264 log point (SE $=$ 0.033) decrease in cigarette consumption. This reflects a 30 percent decrease. In Columns 2–5, we separately add each of the nonprice factors as controls. We find that the variation in antismoking appropriations, newspaper discussion of cigarettes,[20] and the place-based legal restrictions index all have strongly statistically significant associations with a decrease in cigarette consumption. We find qualitatively similar — but not statistically significant — results for tobacco industry political donations. Column 6 controls for all of the nonprice factors. We again find strongly significant associations with all nonprice factors except tobacco industry political donations.

Contrasting the estimated coefficient on the state tax across columns illustrates a striking consequence of these findings. Across Columns 2–5, we find that the inclusion of individual controls consistently reduces the estimated coefficient on state tax (although in Columns 2 and 5 the 95 percent confidence intervals still include the Column 1 estimate). As seen in Column 6, including all of these

[19] Despite this focus, we note that we also find significant coefficients on nonprice factors in control states (see Table A1).

[20] Headlines containing the word “cigarette” may matter through two main channels. First, the headline could contain information about nonprice factors, including the health consequences of smoking and social norms. Second, the headline could contain information about the expected or actual future increase in the price of cigarettes. To explore this distinction, we separate the cigarette headlines into those that contain the word tax and those that do not contain the word tax. Table A2 shows that the estimated effect in Table 2 is driven exclusively by cigarette headlines that do not contain the word “tax,” providing no evidence that the headlines are changing behavior by providing information more relevant to taxes than to smoking itself.

Table 2

Within-Tax-Change-Event Predictors of Cigarette Consumption

	ln(Cigarettes)					
	(1)	(2)	(3)	(4)	(5)	(6)
State tax	−0.264***	−0.257***	−0.196***	−0.203***	−0.227***	−0.127***
	(0.033)	(0.034)	(0.032)	(0.026)	(0.031)	(0.023)
Tobacco industry political donations per 1,000 citizens (cumulative)		−0.070				−0.022
		(0.089)				(0.054)
Antismoking appropriations per 1,000 citizens (cumulative)			−0.005***			−0.004**
			(0.002)			(0.001)
News headlines per journal (cumulative)				−0.005***		−0.005***
				(0.001)		(0.001)
Place-based legal restrictions index					−0.152***	−0.087**
					(0.031)	(0.034)
State-event fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,642	9,642	9,642	9,642	9,642	9,642
Events	150	150	150	150	150	150

Note: This table reports the estimated coefficients from Equation (2). Standard errors are in parentheses and are corrected using multidimensional clustering that allows for correlation within event, within state, and within year-month.

* $p < 0.10$.

** $p < 0.05$.

*** $p < 0.01$.

controls reduces the estimated responsivity to the tax by 52 percent, from −0.264 (SE = 0.033) with no controls to −0.127 (SE = 0.023) with all the controls. The difference between these estimates is strongly statistically significant ($p < 0.001$).

These analyses suggest that a substantial portion of the change in demand that could be attributed to tax changes per se is better attributed to typically unobserved underlying trends in nonprice factors. In the presence of this type of problem, a common solution is to examine particularly narrow windows of time around an event, relying on the logic that a rapid response must be better attributable to the candidate event than any cooccurring trends over time. To assess such a possibility, and to generally investigate the robustness of our findings, we reestimate our model with event windows of varying lengths. Figure 4 reports the coefficients on state tax estimated with and without controls for nonprice factors (as in Columns 6 and 1 of Table 2), using event windows ranging in size from 12 months of pre- and postevent data to 60 months of pre- and postevent data. Estimated price effects are systematically larger for longer time windows, reflecting the fact that over longer windows the underlying trends in nonprice factors have greater confounding effects. However, for all event windows widths presented, the inclusion of nonprice controls substantially reduces the estimated effect. Using the narrowest event window, the estimate decreases in magnitude from −0.124 without the nonprice controls to

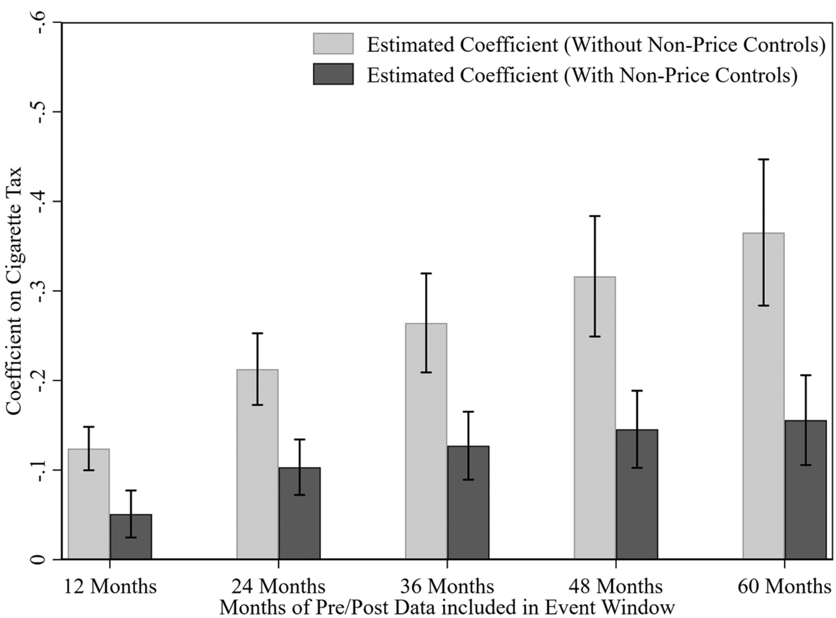

Figure 4. Estimated price responsivity with and without nonprice controls. This figure reports the responsivity of demand to the state tax level from estimating Equation (2). Across columns, we vary the width of the event window. Contrasting the light and dark gray bars illustrates the degree to which price responsivity declines when the nonprice factors are included as control variables.

−0.051 with nonprice controls — a 59 percent decline. In summary, much of the behavioral response around a tax change is explained by variation in nonprice factors regardless of whether short- or long-term pre- or postcomparisons are considered.

Table 2 and Figure 4 clearly demonstrate that controlling for our four nonprice factors substantially attenuates the relationship between cigarette consumption and state taxes. Recall that we have argued that these four factors, while important, are only imperfect proxies for the full class nonprice factors potentially influencing demand. To generate approximate bounds on how much attenuation would occur if all nonprice factors were fully controlled, we apply the approach of Oster (2019). Conceptually, Oster's approach considers a case where the estimation of a treatment effect is confounded by both observable and unobservable controls. Given knowledge of the stability of the treatment effect coefficient as observed controls are added, and given assumptions on the comparative selection occurring from observed and unobserved controls and the total predictive power (R^2) that would be achieved if all controls were observed, Oster's approach yields an estimate of the fully unconfounded treatment effect.

Even with conservative assumptions on the predictive importance of unobserved nonprice factors, Oster's approach suggests substantially more attenuation could be expected. Comparing the regressions from Columns 1 and 6 of Table 2, the inclusion of the four nonprice factors increased the within-event R^2s from 0.217 to 0.383 (i.e., including our four nonprice factors generated a 76 percent increase in explained variation from the baseline specification). For this exercise, we assume that the inclusion of additional unobserved nonprice factors could increase the explained variation by no more than an additional 10 percent (for a within-event R^2 of no more than 0.421). Even with this relatively conservative assumption on the importance of unobserved variables, the corrected magnitude of the coefficient on state tax could be as low as −0.060.[21] This estimate reflects a 77 percent decrease from the estimate without controls for nonprice factors of −0.264. Of course, if the upper bound on the potential explanatory power of remaining unobserved variables is further relaxed, the lower bound on the state tax coefficient would be even smaller.

In the Online Appendix, we provide additional analyses and discussion relevant for assessing the robustness of the results of this section. In particular, we assess the importance of modeling effects as cumulative, the role of linearity assumptions, the consequences of reweighting observations based on the number of events in a state, and issues related to the selection inherent in a sample of pregnant women having live births. While all these issues are conceptually substantive, we present analyses suggesting that they do not confound the qualitative inferences that we draw in this paper. We also examine how our nonprice factors individually contribute to reducing estimated responsivity. These results provide additional nuance to our discussion,

[21] Bootstrapped standard error on bound: 0.019, based on 1,000 bootstrap iterations.

but consistently support our overarching point that nonprice factors have a large influence in these analyses.

The results of this section may be interpreted in two ways. First, and most directly, each regression presented earlier provides an estimate of behavioral responsivity to the tax (formally, a semielasticity) that is identified from within-state variation in a narrow time window and is valid under the assumption that the included nonprice factors are sufficient to completely control for all relevant omitted variables. As we have cautioned, we believe that we likely are still only imperfectly measuring the necessary controls, so we view these estimates as revealing the (imperfect) inference a researcher would make when assuming a given set of observed controls was enough to proceed. Our results imply that relying on estimates without nonprice factor controls would overstate the responsivity of cigarette consumption to taxes.[22]

Second, and more technically, these analyses reveal that the within-treated-state association between cigarette consumption and taxes is significantly confounded. This presents a challenge not only for our specific regressions but also for a broad class of quasi-experimental analyses. Common identification strategies operate by contrasting the within-treated-state association with that in counterfactual predicted associations in which the relevant confounding is assumed to be controlled in some way. The difference is then interpreted as the causal effect of the tax itself. Our results illustrate that the most natural means of controlling for these confounding factors will be highly imperfect. For example, the fact that nonprice factors evolve differently in states with and without tax changes, combined with the finding that controlling for the evolution of nonprice factors strongly influences the association between cigarette consumption and state taxes, suggests that difference-in-difference designs can at best imperfectly account for these concerns. Moreover, the fact that substantial within-treated-state confounding exists even in a narrow window around a tax change suggests that controlling for all nonprice factors with a general, state-specific time trend will be insufficient unless the time trend is sufficiently flexible to deviate from its broad trend is significant ways in a narrow period of time around the event. Convincingly establishing causation in the presence of such loosely specified state-specific time trends is extremely challenging and, in practice, is only successful when

22 To benchmark our quantitative effects relative to existing literature, it is helpful to consider a log-log specification (i.e., an elasticity) rather than our default log-linear specification (i.e., a semielasticity). The log-log version of Table 2's analysis yields an estimated tax elasticity of -0.566 (SE $=$ 0.068) without nonprice-factor controls and of -0.286 (SE $=$ 0.049) with all four nonprice-factor controls. The uncontrolled estimate falls within the typical range of adult price elasticity estimates in this literature, and our fully controlled estimates reflect a greater degree of inelasticity than most estimates. For example, Chaloupka and Warner (2000) report a consensus price elasticity range of -0.4 to -0.7, and Gallet and List (2003) report a mean price elasticity of -0.48 from a meta-analysis of 523 studies. These articles (and more) are discussed in the recent survey of DeCicca, Kenkel, and Lovenheim, (2022). Related to our findings, DeCicca, Kenkel, and Lovenheim (2022) additionally document that lower elasticities are estimated in the relatively uncommon recent studies that attempt to control for antismoking sentiments.

there is clear evidence of discontinuities in behavior at the time of an event. In the following section, we present additional analyses geared toward examining evidence for such discontinuities.

C. Testing Additional Predictions of the Influence of Nonprice Factors

The results above suggest that much of the evolution in cigarette consumption that occurs around a cigarette tax change might be explained by the coevolution of nonprice factors. This finding has strong implications for the predicted timing of the decrease in consumption. Because the evolution of nonprice factors begins well before the tax is enacted, the decrease in consumption that occurs around a tax change would be expected to begin well in advance of when a tax change becomes relevant for prices.

To test the prediction, Figure 5 plots the time path of cigarette consumption in treated states compared with control states. Formally, we estimate Equation (1) using the natural logarithm of cigarette consumption as the dependent variable and using one-month bins. As in the analysis of Subsection IV.A, we allow the period

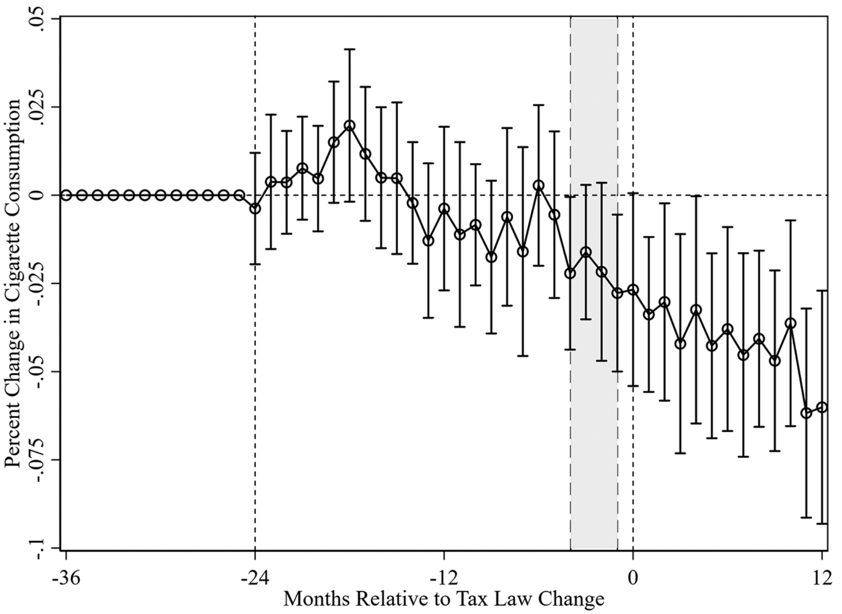

Figure 5. Time path of consumption around tax changes. This figure reports difference-in-differences estimates illustrating the decline in cigarette consumption in states experiencing a tax change. The reported coefficients come from estimating Equation (1) with the natural log of the cigarette consumption as the dependent variable and with time binned at the month level. Capped lines indicate 90 percent confidence intervals. The shaded region indicates the 25th–75th percentile range of dates at which votes on tax law changes occurred.

Table 3

Within-Tax-Change-Event Predictors of Anticipatory Cigarette Consumption Response

	ln(Cigarettes)					
	(1)	(2)	(3)	(4)	(5)	(6)
Effective state tax	−0.266***	−0.259***	−0.198***	−0.205***	−0.229***	−0.128***
	(0.033)	(0.034)	(0.033)	(0.026)	(0.032)	(0.023)
Size of enacted tax change (prior to being effective)	−0.081***	−0.076***	−0.069***	−0.050**	−0.062***	−0.030
	(0.021)	(0.023)	(0.022)	(0.020)	(0.023)	(0.021)
Tobacco industry political donations per 1,000 citizens (cumulative)		−0.066				−0.021
		(0.089)				(0.054)
Antismoking appropriations per 1,000 citizens (cumulative)			−0.005***			−0.004**
			(0.002)			(0.001)
News headlines per journal (cumulative)				−0.005***		−0.005***
				(0.001)		(0.001)
Place-based legal restrictions index					−0.150***	−0.086**
					(0.031)	(0.034)
State-event fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,642	9,642	9,642	9,642	9,642	9,642
Events	150	150	150	150	150	150

Note: This table reports the estimated coefficients from Equation (2) with the inclusion of a variable for the tax rate that has been enacted but is not yet effective, labeled "Size of Enacted Tax Change (Prior to Being Effective)." Standard errors are in parentheses and are corrected using multidimensional clustering that allows for correlation within event, within state, and within year-month.

* $p < 0.10$.

** $p < 0.05$.

*** $p < 0.01$.

of time 36 months before the tax change to 25 months before the tax change to serve as the baseline period, and then we examine the evolution of our dependent variable in the following 24 months prior to the tax change and the 12 months following the tax change. The figure reports our estimates of the differential change occurring in treated states in the event time.

Standard models of price sensitivity predict a discontinuous decrease in demand occurring precisely at the moment the tax changes. Figure 5 provides no clear evidence of such a discontinuity, and instead suggests that the decrease in demand occurs as a gradual decrease over a number of years.

Models of rational addiction predict some degree of anticipatory decrease in demand arising from a known future increase in prices. Prior work documents decrease in demand occurring before the tax is enacted but after a vote has occurred and use this difference to support the existence of rational addiction (see, e.g., Gruber and Kőszegi, 2001). Note that the anticipatory effect we document here is occurring substantially earlier (the shaded region of Figure 5 illustrates the interquartile range of the enactment of tax changes in our data). Of course, a smoker who is aware of upcoming votes might hold a rational expectation of some future price increase even before the vote occurs, which could contribute to the earlier decrease in consumption that we observe.

Building on these observations, we investigate how accommodating the regression framework of Gruber and Kőszegi (2001) affects our results. The key innovation of their approach is to regress cigarette consumption on not only the tax rate that is in effect but also on the tax rate that has been enacted but is not yet effective. The coefficient on this latter term is interpreted as capturing anticipatory response to the future tax increase as in the rational-addiction model.

Table 3 reports results for regressions like those in Table 2 but includes a variable for the enacted but not yet effective tax. In the period between when a tax change is enacted and when it becomes effective, this variable takes the value of the size of the tax change (measured in dollars). At all other times, this variable takes the value of zero. Column 1 reproduces the qualitative result of Gruber and Kőszegi (2001), showing that 30 percent of the ultimate decrease in consumption that is attributed to the tax occurs in the window after the tax is enacted but before the tax goes into effect. Upon examining Columns 2–5, however, we see that even with the inclusion of this variable, the relevance of the nonprice factors persists. In Column 6, which reproduces our main specification with all nonprice controls included, we continue to find that the estimate on tax is reduced by more than half relative to the baseline specification in Column 1. Furthermore, the coefficient on the Gruber and Kőszegi measure of anticipation of the price change becomes smaller and is not statistically significant at conventional levels. This suggests that the strength of the statistical evidence in support of Gruber and Kőszegi's account is meaningfully weakened when nonprice factors are accommodated. We interpret these findings to illustrate the difficulty of causally attributing declines in consumption around tax changes to (even anticipatory) price events, but we emphasize that the results do not reject meaningful anticipatory effects.

V. CONCLUSION

In the standard economic framework, sin taxes dissuade behavior through price effects. All else being equal, a raise in taxes raises market prices, which in turn reduces demand. While this logic is correct and compelling, in this article, we document that all else is not held equal in the course of a tax law change and that the behavioral response observed around a tax change might largely be attributable to nonprice factors.

Within the sin-tax literature, we highlight the importance of our findings in three ongoing debates. First, the results inform the ongoing policy discussion about the expected effects of sin taxes. States have long employed taxes on items like cigarettes and alcohol, but the recent surge of interest in taxing soda and sugary beverages has renewed conceptual interest in this domain. As we proceed with a wave of attempts to impose new such taxes, it is important to understand the key determinants of successfully dissuasive tax changes. Our results help inform the channels through which these legal changes can achieve their goal and illustrate the potential for nonprice channels to drive observed effects. Given that both price-based and nonprice-based channels are important, and given that these channels are commonly activated in tandem, understanding the optimal mix of such policies remains an important open question for future research.

Second, the results are relevant to the literature aimed at assessing the evidence of "rational addiction" in the spirit of Becker and Murphy (1988). Classic tests of this theory, such as Gruber and Kőszegi (2001), present evidence that cigarette consumption decreases in advance of the tax change but after the tax has been enacted but is not yet effective. This is interpreted as evidence that forward-looking smokers reduce their degree of addiction in anticipation of cigarettes becoming more expensive. Our results suggest that a variety of other factors can contribute to decreases in consumption preceding the tax change. Our conceptual framework does not rely on the absence of any rational-addiction motives — indeed, the provision of information that we document could interact with these motives in a manner that helps reduce demand. However, our results suggest that interpreting anticipatory decrease in consumption as clear evidence of rational addiction could be done only under the stringent scenario in which all nonprice factors are fully controlled for in empirical analysis.

Third, the results are relevant to the estimation and interpretation of sin-tax elasticities. We document that the inclusion of our candidate nonprice factors as controls reduces the estimated degree of price sensitivity by roughly half. Furthermore, we emphasize that we have only a partial list of nonprice factors, and our available measures are imperfect proxies for the underlying constructs of interest. As a result, the reduction in the estimated degree of price sensitivity is likely an underestimate. These findings build on prior research that demonstrates that controls for survey measures of antismoking sentiments can reduce estimated price elasticities (e.g., DeCicca et al., 2008). Relative to these past findings, we contribute by directly studying a variety of the factors that we believe shape antismoking sentiments, by demonstrating that these factors evolve rapidly in the course of a tax change, and by documenting that this rapid evolution may account for a large portion of the behavioral response occurring in the vicinity of tax changes.

Beyond the context of sin taxation, the results align with recent findings in several disparate domains demonstrating that more than financial incentives change in the course of a policy change. Gneezy and Rustichini (2000) find that introducing financial penalties for late pick-ups from childcare led to increased late pick-up, which is argued to be driven by the impact of the fine on perceptions of the social consequences of the behavior. Abouk and Adams (2013) find that the decrease of accidents after texting bans is better explained by the bans' announcements than the bans themselves. Malani and Reif (2015) find that physician labor supply anticipatory reactions to tort reform significantly affects the estimated effect of the reform on behavior. Richwine, Dor, and Moghtaderi (2019) find an increase in both public discourse about the benefits of vaccinations and vaccinations themselves in the lead-up to a law that removed all nonmedical exemptions for them in schools. Taylor et al. (2019) find that the decline in soda consumption occurring during the Berkeley Soda Tax rollout can be largely explained by impacts of election and social media coverage. Finally, Pratt (2019) finds that the impact of fines associated with water rationing can be partially explained by their impact on the perceived importance of conservation.

We empirically contribute to this literature by showing that the types of concerns raised in the previous work are likely to be first order in models of optimal sin taxation. Conceptually, this diverse body of related examples can be unified through the legal literature concerning how laws work and why they are enacted. Legal scholars have long appreciated that laws do more than change financial incentives and can change behavior through what is commonly referred to as expressive effects. Our work supports and expands on the concept of expressive effects that incorporates responses to the ideas communicated by all the activities and debate that surround and accompany the process of legal change. Despite the prevalence of the view that law can work by expressing value, this article presents some of the first empirical evidence that begins to separate price effects from expressive effects. Our results suggests that, at least in the context of cigarette taxes, the manner in which tax laws operate through nonprice channels should be a bigger consideration for both researchers and policy makers.

ACKNOWLEDGMENTS

We thank Douglas Baird, Adam Chilton, Erin Delaney, Gerrit De Geest, Dhammika Dharmapala, Michael Frakes, Ezra Friedman, Daniel Hemel, Jim Hines, William Hubbard, Don Kenkel, Genevieve Lakier, Sarah Lawsky, Ben Lockwood, Anup Malani, Jonathan Masur, Richard McAdams, Adam Rosenzweig, Sarath

Sanga, Max Schanzenbach, Joel Slemrod, Kate Smith, Dmitry Taubinsky, and seminar audiences at the American Economic Association Annual Meeting, the American Law and Economics Association Annual Meeting, the Conference on Empirical Legal Studies, the Junior Tax Scholars Workshop, George Mason University Antonin Scalia Law School, University of California Berkeley Haas School of Business, University of California Berkeley School of Law, University of California Hastings College of Law, University of Chicago Law School, University of Southern California School of Law Center for Law and Social Science, Washington University School of Law, and Wharton for helpful comments. We thank Andrew Joung for excellent research assistance and the Wharton Dean's Research Fund for financial support.

DISCLOSURES

The authors have no financial conflicts to disclose.

REFERENCES

Abouk, R., and S. Adams, 2013. "Texting Bans and Fatal Accidents on Roadways: Do They Work? Or Do Drivers Just React to Announcements of Bans?" *American Economic Journal: Applied Economics* 5 (2), 179–199.

Allcott, Hunt, Benjamin Lockwood, and Dmitry Taubinsky, 2019. "Regressive Sin Taxes, with an Application to the Optimal Soda Tax." *Quarterly Journal of Economics* 134 (3), 1557–1626.

Becker, Gary, and Kevin Murphy, 1988. "A Theory of Rational Addiction." *Journal of Political Economy* 96 (4), 675–700.

Beleche, Trinidad, Nellie Lew, Rosemarie L. Summers, and J. Laron Kirby, 2018. "Are Graphic Warning Labels Stopping Millions of Smokers? A Comment on Huang, Chaloupka, and Fong." *Econ Journal Watch* 15 (2), 129–157.

Binder, M., and M. H. Pesaran, 2001. "Life-Cycle Consumption under Social Interactions." *Journal of Economic Dynamics and Control* 25 (1–2), 35–83.

Blanchflower, D. G., B. Van Landeghem, and A. J. Oswald, 2009. "Imitative Obesity and Relative Utility." *Journal of the European Economic Association* 7 (2–3), 528–538.

Bonica, Adam, 2016. *Database on Ideology, Money in Politics, and Elections: Public Version 1.0 (Computer File)*. Stanford University Libraries, Stanford, CA, https://data.stanford.edu/dime.

Brock, William A., and Steven N. Durlauf, 2001. "Discrete Choice with Social Interactions." *Review of Economic Studies* 68 (2), 235–260.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134 (3), 1405–1454.

Centers for Disease Control and Prevention, 2018. "State Tobacco Activities Tracking and Evaluation System." https://chronicdata.cdc.gov/Legislation/CDC-STATE-System-Tobacco-Legislation-Smokefree-Ind/32fd-hyzc.

Centers for Disease Control's Health Communication Interventions, 2015. "Best Practices Program Components from the Centers for Disease Control and Prevention." https://chronicdata.cdc.gov/Funding/University-of-Illinois-at-Chicago-Health-Policy-Ce/vw7y-v3uk.

Chaloupka, Frank J., and Kenneth E. Warner, 2000. "The Economics of Smoking." In Culyer, Anthony J., and Joseph P. Newhouse (eds.), *Handbook of Health Economics*, Vol. 1, Part B, 1539–1627. North Holland, New York.

Chiou, Lesley, and Erich Muehlegger, 2008. "Crossing the Line: Direct Estimation of Cross-Border Cigarette Sales and the Effect on Tax Revenue." *B.E. Journal of Economic Analysis and Policy* 8 (1), 1–41.

Chiou, Lesley, and Erich Muehlegger, 2014. "Consumer Response to Cigarette Excise Tax Changes." *National Tax Journal* 67 (3), 621–650.

Correia, Sergio, 2016. "A Feasible Estimator for Linear Models with Multi-Way Fixed Effects." Working Paper. Duke University, Durham, NC, https://scorreia.com/research/hdfe.pdf.

DeCicca, Philip, Donald Kenkel, and Feng Liu, 2013a. "Excise Tax Avoidance: The Case of State Cigarette Taxes." *Journal of Health Economics* 32 (6), 1130–1141.

DeCicca, Philip, Donald Kenkel, and Feng Liu, 2013b. "Who Pays Cigarette Taxes? The Impact of Consumer Price Search." *Review of Economics and Statistics* 95 (2), 516–529.

DeCicca, Philip, Donald Kenkel, and Michael F. Lovenheim, 2022. "The Economics of Tobacco Regulation: A Comprehensive Review." *Journal of Economic Literature* 60 (3), 883–970.

DeCicca, Philip, Donald S. Kenkel, Alan D. Mathios, and Justine Shin, 2008. "Youth Smoking, Cigarette Prices and Anti-smoking Sentiment." *Health Economics* 17 (6), 733–749.

Deshpande, Manasi, and Yue Li, 2019. "Who Is Screened Out? Application Costs and the Targeting of Disability Programs." *American Economic Journal: Economic Policy* 11 (4), 213–248.

Dwenger, Nadja, Henrik Kleven, Imran Rasul, and Johannes Rincke, 2016. "Extrinsic and Intrinsic Motivations for Tax Compliance: Evidence from a Field Experiment in Germany." *American Economic Journal: Economic Policy* 8 (3), 203–232.

Evans, William N., Jeanne S. Ringel, and Diana Stech, 1999. "Tobacco Taxes and Public Policy to Discourage Smoking." *Tax Policy and the Economy* 13, 1–56.

Fabbri, Marco, and Sven Hoeppner, 2018. "Compliance Externalities and the Role-Model Effect on Law Abidance: Field and Survey Experimental Evidence." *Journal of Empirical Legal Studies* 15 (3), 539–562.

Farhi, Emmanuel, and Xavier Gabaix, 2017. "Optimal Taxation with Behavioral Agents." NBER Working Paper No. 21524. National Bureau of Economic Research, Cambridge, MA.

Funk, Patricia, 2007. "Is There An Expressive Function of Law? An Empirical Analysis of Voting Laws with Symbolic Fines." *American Law and Economics Review* 9 (1), 135–159.

Gallet, Craig A., and John A. List, 2003. "Cigarette Demand: A Meta-Analysis of Elasticities." *Health Economics* 12 (10), 821–835.

Glaeser, Edward L., and José A. Scheinkman, 2002. "Nonmarket Interactions." In Dewatripoint, Mathias, Lars P. Hansen, and Stephen J. Turnovsky (eds.), *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, 339–369. Cambridge University Press, Cambridge.

Gneezy, Uri, and Aldo Rustichini, 2000. "A Fine Is a Price." *Journal of Legal Studies* 29 (1), 1–17.

Gruber, Jon, and Botond Kőszegi, 2001. "Is Addiction 'Rational'? Theory and Evidence." *Quarterly Journal of Economics* 116 (4), 1261–1303.

Hamilton, William L., Lois Biener, and Robert T. Brennan, 2008. "Do Local Tobacco Regulations Influence Perceived Smoking Norms? Evidence from Adult and Youth Surveys in Massachusetts." *Health Education Research* 23 (4), 709–722.

Harberger, Arnold, 1964. "The Measurement of Waste." *American Economic Review* 54 (3), 58–76.

Kessler, Judd B., and Michael I. Norton, 2016. "Tax Aversion in Labor Supply." *Journal of Economic Behavior and Organization* 124, 15–28.

Lovenheim, Michael, 2008. "How Far to the Border?: The Extent and Impact of Cross-Border Casual Cigarette Smuggling." *National Tax Journal* 61 (1), 7–33.

Malani, Anup, and Julian Reif, 2015. "Interpreting Pre-Trends as Anticipation: Impact on Estimated Treatment Effects from Tort Reform." *Journal of Public Economics* 124, 1–17.

McAdams, Richard H., 2015. *The Expressive Powers of Law: Theories and Limits*. Harvard University Press, Cambridge, MA.

McAdams, Richard H., and Janice Nadler, 2005. "Testing the Focal Point Theory of Legal Compliance: The Effect of Third–Party Expression in an Experimental Hawk/Dove Game." *Journal of Empirical Legal Studies* 2 (1), 87–123.

O'Donoghue, Ted, and Matthew Rabin, 2003. "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes." *American Economic Review Papers and Proceedings* 93 (2), 186–191.

O'Donoghue, Ted, and Matthew Rabin, 2005. "Optimal Taxes for Sin Goods." *Swedish Economic Policy Review* 12 (2), 7–39.

O'Donoghue, Ted, and Matthew Rabin, 2006. "Optimal Sin Taxes." *Journal of Public Economics* 90 (10), 1825–1849.

Oster, Emily, 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business and Economic Statistics* 37 (2), 187–204.

Pigou, Arthur, 1920. *The Economics of Welfare*. Macmillan, London.

Pratt, Bryan, 2019. "A Fine is More Than a Price: Evidence from Drought Restrictions." *Mimeograph*.

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

Reif, Julian, 2019. "A Model of Addiction and Social Interactions." *Economic Inquiry* 57, 759–773.

Richwine, Chelsea J., Avi Dor, and Ali Moghtaderi, 2019. "Do Stricter Immunization Laws Improve Coverage? Evidence from the Repeal of Non-medical Exemptions for School Mandated Vaccines." NBER Working Paper No. 25847. National Bureau of Economic Research, Cambridge, MA.

Riley, Kristen E., Michael R. Ulrich, Heidi A. Hamann, and Jamie S. Ostroff, 2017. “Decreasing Smoking but Increasing Stigma? Anti-tobacco Campaigns, Public Health, and Cancer Care.” *AMA Journal of Ethics* 19 (5), 475–485.

Roberto, Christina, Diandra Wong, Aviva Musicus, and David Hammond, 2016. “The Influence of Sugar-Sweetened Beverage Health Warning Labels on Parents’ Choices.” *Pediatrics* 137 (2), 1–10.

Smulian, John C., Cande V. Ananth, Maryellen L. Hanley, Robert A. Knuppel, Judy Donlen, and Lakota Kruse, 2001. “New Jersey’s Electronic Birth Certificate Program: Variations in Data Sources.” *American Journal of Public Health* 91 (5), 814–816.

Stuber, Jennifer, Sandro Galea, and Bruce G. Link, 2009. “Stigma and Smoking: The Consequences of Our Good Intentions.” *Social Service Review* 83 (4), 585–609.

Taylor, Rebecca, Scott Kaplan, Sofia Villas-Boas, and Kevin Jung, 2019. “Soda Wars: The Effect of a Soda Tax Election on University Beverage Sales.” *Economic Inquiry* 57, 1480–1496.

Wittlin, Maggie, 2011. “Buckling under Pressure: An Empirical Test of the Expressive Effects of Law.” *Yale Journal on Regulation* 28 (2), 419–469.