NBER WORKING PAPER SERIES

OPTIMAL REGULATION OF E-CIGARETTES: THEORY AND EVIDENCE

Hunt Allcott
Charlie Rafkin

Working Paper 27000
http://www.nber.org/papers/w27000

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020

We thank Alberto Abadie, Josh Angrist, Noel Brewer, Chad Cotti, Charles Courtemanche, Brigham Frandsen, Don Kenkel, Jessica Pepper, Mike Pesko, Kurt Ribisl, Frank Schilbach, Andrew Seidenberg, Jann Spiess, Tevi Troy, and seminar participants at ITAM, Microsoft Research, and MIT for helpful feedback. We thank Raj Bhargava, Ahmad Rahman, and Aakaash Rao for exceptional research assistance. We are grateful to the Sloan Foundation for financial support. This material includes work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 (Rafkin). Replication files and survey instruments are available from https://sites.google.com/site/allcott/research. This paper represents the researchers' own analyses calculated based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen nor the National Bureau of Economic Research. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

Optimal Regulation of E-cigarettes: Theory and Evidence
Hunt Allcott and Charlie Rafkin
NBER Working Paper No. 27000
April 2020
JEL No. D12,D18,D61,H21,H23,I12,I18

ABSTRACT

We model optimal e-cigarette regulation and estimate key sufficient statistics. Using tax changes and scanner data, we estimate relatively elastic demand and limited substitution between e-cigarettes and combustible cigarettes. In sample surveys, historical smoking declines for high- and low-vaping demographics were unchanged after e-cigarettes were introduced; this demographic shift-share identification also suggests limited substitution. We field a new survey of experts, who report that vaping is almost as harmful as smoking cigarettes. In our model, these results imply that current e-cigarette taxes are far below the social optimum, but Monte Carlo simulations highlight substantial uncertainty.

Hunt Allcott
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
hunt.allcott@nyu.edu

Charlie Rafkin
MIT Department of Economics
77 Massachusetts Avenue
Cambridge, MA 02139
crafkin@mit.edu

As of 2019, eight million American adults and four million American youth report using e-cigarettes, and many more youth now vape e-cigarettes than smoke traditional combustible cigarettes. There is significant disagreement about whether regulators should encourage or discourage this popular new product. Optimists point out that the widespread adoption suggests that e-cigarettes generate substantial consumer surplus. Furthermore, e-cigarettes can be a useful smoking cessation aid (Hajek et al. 2019), and vaping may be less harmful than smoking cigarettes (National Academy of Sciences 2018). On the other hand, pessimists point out that widespread adoption of an addictive product is not necessarily good for well-being. Furthermore, vaping might be a gateway to smoking for youth, and the exact health effects of vaping are uncertain, as underscored by a recent spate of vaping-related illnesses and deaths (Gotts et al. 2019).

This disagreement has played out in divergent and sometimes conflicting policies. Three-quarters of Americans live in places with no e-cigarette taxes, while the states and local areas that do tax e-cigarettes impose very different rates. Many regulators think of e-cigarettes as a promising harm reduction tool for current smokers (Gottlieb 2018; Zeller 2019), but San Francisco has effectively banned all e-cigarette sales while keeping combustible cigarettes legal.

Is vaping in fact a substitute for smoking cigarettes, or a complement? Is this different for youth versus adults? What is the state of expert knowledge about the relative harms of vaping versus smoking? What is the socially optimal e-cigarette tax rate? Could it be optimal to ban all e-cigarette sales? How certain can we be about any policy prescriptions? This paper lays out a model of optimal e-cigarette regulation and derives equations for optimal taxes and welfare as functions of several sufficient statistics. We then estimate key statistics using an array of empirical data and propose answers to the above questions.

Our theoretical model extends the sufficient statistic approach to behavioral public economics (Chetty, Looney and Kroft 2009; Mullainathan, Schwartzstein and Congdon 2012; Allcott and Taubinsky 2015; Bernheim and Taubinsky 2018; Farhi and Gabaix 2020) to a dynamic setting appropriate for studying addictive goods. We model heterogeneous consumers who consume a numeraire good plus two habit-forming goods (cigarettes and e-cigarettes) that impose internalities and externalities. By "internalities," we mean that the social planner believes that consumers' choices do not maximize their own long-run utility, perhaps because of present focus, projection bias or related misperceptions of addiction, or biased beliefs about health harms.[1] The social planner can tax or ban either good.

In this framework, the optimal e-cigarette tax depends on three sufficient statistics: the marginal uninternalized harms (externalities and internalities) from vaping, the marginal uninternalized harms from smoking cigarettes, and the extent to which vaping and smoking are complements

[1] For more discussion and evidence on internalities related to smoking and vaping, see Viscusi (1990; 2016), Gruber and Koszegi (2001; 2004), Gruber and Mullainathan (2005), Chaloupka et al. (2015), Ashley, Nardinelli and Lavaty (2015), Cutler et al. (2015; 2016), Jin et al. (2015), DeCicca et al. (2017), Kenkel et al. (2019), Levy, Norton and Smith (2018), Chaloupka, Levy and White (2019), and DeCicca, Kenkel and Lovenheim (2020).

or substitutes. The welfare effect of banning e-cigarettes compared to keeping taxes at current levels depends on those same statistics plus the perceived consumer surplus loss as revealed by the e-cigarette demand curve. Optimally set taxes are always preferred to a ban in our model, but a ban may increase welfare relative to the status quo if tax rates are constrained by political issues, tax evasion, or other factors. Furthermore, a type-specific ban (for example, a youth sales ban) may be optimal given that uninternalized harms vary across types and type-specific taxes are hard to implement.

To estimate e-cigarette demand, we use Nielsen scanner data on e-cigarette sales at 27,000 stores across the country from 2013–2017. To identify the price elasticity, we exploit changes in state and local e-cigarette taxes. Before the tax changes, there is no trend in retail prices or quantities sold. After the tax changes, tax-inclusive retail prices rise and persistently, and quantities sold drop. Our primary estimate suggests an own-price elasticity of about -1.32. We also estimate the elasticity of substitution between e-cigarettes and cigarettes using tax changes and sales for both goods. Our primary estimates suggest statistically insignificant substitutability. However, the point estimates vary somewhat across specifications, and aggregate scanner data cannot identify heterogeneous substitution parameters: vaping and smoking could still be substitutes for adults and complements for youth.

We thus turn to a more novel strategy to identify substitution patterns, exploiting the fact that different demographic groups have very different demand for e-cigarettes. Specifically, White people, men, non-college graduates, lower-income people, and younger adults (but older youth) vape more than non-Whites, women, etc. Some of these demand differences may be related to broader preferences for new technologies: we show that the demographics of e-cigarette early adopters—in particular, their age profile—is similar to the demographics of internet and social media early adopters. Between 2004 and 2012, i.e. before e-cigarettes became popular, the demographic groups that would later have higher e-cigarette demand had steady linear declines in cigarette smoking relative to demographics with lower latent demand. If that relative decline accelerated after e-cigarettes became popular, this would suggest that vaping caused smoking to decrease, and thus that e-cigarettes are substitutes for combustible cigarettes. On the other hand, if that relative decline slowed, this would suggest that vaping caused more smoking, and thus that e-cigarettes are a gateway to combustible cigarettes.

This approach is a cousin of the "shift-share" identification strategy popularized by Bartik (1991): we primarily exploit cross-sectional variation in demand across demographics with the time-series growth in e-cigarette use. The identifying assumption is that any changes in relative smoking trends for high- versus low-vaping demographics were caused by the introduction of e-cigarettes. In support of this assumption, we find that smoking decreases were close to linear in the years before e-cigarettes were introduced and that the estimates are consistent across different demographics.

We implement this demographic shift-share strategy using data from five large nationally representative surveys comprising 7.4 million observations collected over 2004–2018: the Behavioral Risk Factor Surveillance Survey, the National Health Interview Survey, the National Survey of Drug Use and Health, Monitoring the Future, and the National Youth Tobacco Survey. Our estimates are consistent with our earlier estimates identified from tax changes in suggesting that on average, vaping is not a significant complement or substitute for smoking. Our confidence intervals rule out that the introduction of e-cigarettes affected the 2004–2018 smoking decrease by more than 5 to 11 percent in either direction. To believe that e-cigarettes increased or decreased smoking by more than that, one would have to think that high-vaping demographics (young adults, White people, men, etc.) coincidentally all had unpredicted decreases or increases in cigarette demand over the past six years that exactly offset the alleged effects of their vaping.

There is great uncertainty about the uninternalized harms from vaping, and research is evolving rapidly. To aggregate the state of knowledge about the harms from e-cigarettes, we surveyed the corresponding authors of papers on the health impacts of e-cigarettes in the landmark National Academy of Sciences (2018) study as well as economists who have written on cigarettes or e-cigarettes. The average expert who responded believes that the external and internal harms from vaping are 48 and 101 percent as large as those from smoking cigarettes. There is substantial disagreement: the interquartile ranges of beliefs about these two relative harms are 10 to 75 percent and 27 to 180 percent, respectively. These results paint a very different picture compared to a prominent and controversial early estimate that vaping was only five percent as harmful as smoking cigarettes (Nutt et al. 2014). We estimate the dollar value of the uninternalized harms from vaping by combining the expert survey results with prior estimates of the internalities and externalities from smoking from Gruber and Kőszegi (2001), Sloan et al. (2004), Cutler et al. (2015), and Chaloupka, Levy and White (2019).

Finally, we use our model to evaluate optimal e-cigarette regulation, using Monte Carlo simulations to account for uncertainty. The above three empirical results have clear implications for optimal policy. Relatively elastic demand implies relatively small perceived consumer surplus losses from an e-cigarette ban. Limited substitutability with combustible cigarettes means that optimal e-cigarette policy depends little on the uninternalized distortions from smoking. Large uninternalized harms from vaping increase the optimal tax rate and increase the welfare gains from a ban compared to current tax rates.

The optimal e-cigarette tax is positive in 97 percent of simulations, and it exceeds the average existing e-cigarette tax in about 95 percent of simulations. It is easy to see why in our model, existing e-cigarette taxes are probably too low. Experts in our survey believe that vaping every day is not much less harmful than smoking every day and that vaping is more harmful for youth than adults. At typical usage, one pack of cigarettes is equivalent to somewhat less than 1 milliliter of e-liquid. Thus, the optimal e-cigarette tax per milliliter will roughly equal the per-pack uninternalized harm from adult smoking. We cite estimates of the adult smoking uninternalized harm ranging from \$5 per pack to over \$80 per pack; relying on Cutler et al. (2015) and Sloan et al. (2004) suggests \$18. By contrast, the average existing e-cigarette tax is only \$0.89/ml.

Since existing e-cigarette taxes are far below optimal in our model, this raises the possibility that a complete ban would generate higher social welfare than current taxes. Eliminating youth vaping increases welfare in 92 percent of model simulations. Thus, the existing regulations banning sales to minors and banning all sales of flavored e-cigarettes (which are especially appealing to minors) likely increase welfare in our model. More controversially, eliminating adult vaping is preferred to the status quo in 90 percent of simulations. At our mean parameter values, fully banning e-cigarettes increases welfare by \$91 per person per year, or \$25 billion per year over the 279 million people aged 12 and older nationwide.

The key caveat to these optimal tax and welfare calculations is that they hinge on assumptions about the uninternalized harms from vaping, and our experts are more pessimistic than prior literature might have suggested. For example, Viscusi (2016) argues that vaping could be at least 100 times safer than smoking and finds that people overestimate vaping risks relative to that benchmark, a bias that would cause people to vape too little. We show that when our model is calibrated with risk misperceptions suggested by Viscusi (2016), it is optimal to heavily subsidize e-cigarettes instead of taxing or banning them. This calculation underscores that our theory, empirical estimates, and policy analyses can be informative even for readers who disagree with our experts about uninternalized harms.

There are a number of additional important caveats. First, the Nielsen scanner data cover only about 2.5 percent of e-cigarette retail, and our price elasticity estimate is biased if this is an unrepresentative sample. Second, because we identify e-cigarette demand off of relatively limited price variation, we must make strong functional form assumptions to estimate inframarginal demand and perceived consumer surplus; this is a standard problem when analyzing the welfare effects of bans or new products (e.g. Hausman 1996). Third, our substitution estimates are identified for a time horizon of several years; we do not yet know if youth vapers will transition to combustible cigarettes later in life or if adult smokers need more time to substitute to e-cigarettes. Fourth, the key parameters may change in the future for any number of reasons, including the coronavirus pandemic and the recent ban on flavored e-cigarettes.

Our work builds on a growing literature on e-cigarettes. Our primary contribution is to provide a framework for modeling optimal policy combined with new estimates of the key empirical parameters. A related paper by Kenkel et al. (2019) presents survey data suggesting that behavioral biases reduce vaping and carries out simulations showing that such behavioral biases against vaping imply that taxing or banning e-cigarettes reduces welfare. As described above, our model makes similar predictions in the presence of behavioral bias against vaping, but our expert survey responses point in the opposite direction.

To our knowledge, our September 2019 working paper was the first estimate of the aggregate price elasticity of e-cigarette demand using tax variation and scanner data (instead of survey data). This distinction may be important: tax changes provide long-run and potentially exogenous variation, and most surveys have imperfect measures of the intensive margin of e-cigarette use. Cotti et al. (2020) released an independent analysis in January 2020, and other papers study the effect of price changes in survey data[2] or use scanner data to estimate different e-cigarette demand parameters.[3]

There is conflicting evidence on whether vaping and smoking are complements or substitutes. A series of papers find that youth who vape are more likely to smoke in the future, even after controlling for observable characteristics that predict both vaping and smoking.[4] Although it is possible that unobserved confounders could cause both smoking and vaping, many public health researchers have taken this as evidence that vaping causes future smoking, and thus that regulating vaping would improve public health.[5] A series of other papers using quasi-experimental strategies have come to the opposite conclusion, finding that vaping and smoking are substitutes. However, there is some disagreement even between papers that use similar identification.[6] Our demographic shift-share approach is novel, and it may help to resolve the disagreement between existing papers.

Our work speaks to three literatures outside of e-cigarettes. First, we extend the behavioral public economics literature on optimal sin taxes (Gruber and Koszegi 2001, 2004; O'Donoghue and Rabin 2006; Allcott and Taubinsky 2015; Bernheim and Taubinsky 2018; Allcott, Lockwood and Taubinsky 2019; Farhi and Gabaix 2020; and others). Second, our demographic shift-share

[2]Pesko and Warman (2017), Pesko et al. (2018), Saffer et al. (2018), and Cantrell et al. (2019) estimate the association between price variation observed in Nielsen scanner data and survey measures of e-cigarette use. Pesko, Courtemanche and Maclean (2019) estimate the effect of cigarette and e-cigarette tax changes on survey measures of e-cigarette use.

[3]Zheng et al. (2017) and Huang et al. (2018) estimate the short run residual demand elasticity faced by particular types of stores, using data at the city-month-store type level. Stoklosa, Drope and Chaloupka (2016) estimate the short-run demand elasticity in the EU using country-by-month data. For our research question, the parameter of interest is the aggregate long-run demand elasticity. Short-run and long-run elasticity may differ due to stockpiling and habit formation, and the residual demand function faced by a set of stores could naturally differ from aggregate demand elasticity as consumers substitute across stores.

[4]See Leventhal et al. (2015), Primack et al. (2015), Watkins, Glantz and Chaffee (2018), Berry et al. (2019), and others, and see Chatterjee et al. (2016) and Soneji et al. (2017) for systematic reviews.

[5]For example, an important review article by Soneji et al. (2017, page 788) concludes that "e-cigarette use was associated with greater risk for subsequent cigarette smoking initiation and past 30-day cigarette smoking. Strong e-cigarette regulation could potentially curb use among youth and possibly limit the future population-level burden of cigarette smoking." Similarly, an earlier review article by Chatterjee et al. (2016, page 1) concludes that "[Electronic cigarettes] are associated with higher incidence of combustible cigarette smoking. Policy makers need to recognize the insidious nature of this campaign by the tobacco industry and design policies to regulate it." The National Academy of Sciences (2018, page 555) study concludes, "the committee considered the overall body of evidence of a causal effect of e-cigarette use on risk of transition from never to ever smoking to be substantial."

[6]Friedman (2015), Pesko and Currie (2019), Pesko, Hughes and Faisal (2016), Cooper and Pesko (2017), Pesko and Warman (2017), Saffer et al. (2018), Saffer et al. (2019), Abouk et al. (2019), Cantrell et al. (2019), Dave, Feng and Pesko (2019), Pesko et al. (2019), and Cotti et al. (2020) find that e-cigarettes and cigarettes are substitutes. Using similar identification (state-level tax variation and bans on e-cigarette sales to minors), however, Abouk and Adams (2017) and Cotti, Nesson and Tefft (2018) find that they are complements.

design is related to Boxell, Gentzkow and Shapiro (2017), who identify the effects of the internet on political polarization by exploiting age differences in internet adoption, and DeCicca et al. (2017), who identify the effects of menthol cigarettes by exploiting racial differences in tastes for menthol. Third, our work is broadly related to studies of the welfare effects of other new products (Trajtenberg 1989; Hausman 1996; Petrin 2002; Nevo 2003; Goolsbee and Petrin 2004; Gentzkow 2007; Aguiar and Waldfogel 2018; and others).

Section 1 lays out the theoretical framework. Sections 2 and 3 present the data and smoking and vaping trends. Sections 4 and 5 present estimates of price elasticity and substitution patterns. Sections 6 and 7 present the expert survey results and optimal policy analysis, and Section 8 concludes.

1 Theoretical Framework

E-cigarette regulation involves setting constant taxes on an addictive good, motivated by both externalities and consumer bias. To match this application, we introduce a dynamic model of consumers who impose externalities and do not necessarily maximize their utility. We then solve for optimal constant tax rates and the welfare effects of banning e-cigarettes compared to keeping taxes at some baseline level. Our model can be thought of as a less parameterized version of the dynamic optimal tax model in Gruber and Koszegi (2001) or a simple dynamic extension of static optimal corrective taxation models such as Diamond (1973), O'Donoghue and Rabin (2006), Allcott and Taubinsky (2015), and Farhi and Gabaix (2020).

1.1 Consumption, Bias, and Welfare

Setup. There are infinite periods indexed by t. There is a numeraire good n and two other goods indexed by j or k: cigarettes c and e-cigarettes e. All goods are produced at constant marginal cost in competitive markets. A social planner sets constant taxes $\boldsymbol{\tau} = \{\tau^c, \tau^e\}$ and maintains a balanced budget in each period using a lump sum transfer T_t. Let $\boldsymbol{p} = \{p^c, p^e\}$ denote the vector of after-tax prices for c and e; n is sold at price 1. While $\boldsymbol{\tau}$ and $\boldsymbol{p}$ might vary in the equations below, let $\tilde{\tau}$ and $\tilde{\boldsymbol{p}}$ denote vectors of baseline taxes and market prices. We write j or k as superscripts to avoid confusion with other subscripts throughout the paper; any time t superscripts are exponents.

Heterogeneous consumers have finite types indexed by θ with measure s_θ and $\sum_\theta s_\theta = 1$. Let $\boldsymbol{q}_t = \{q_t^c, q_t^e\}$ and q_t^n denote possible consumption levels in period t, and let $\boldsymbol{q}_{\theta t} = \{q_{\theta t}^c, q_{\theta t}^e\}$ denote the actual consumption chosen by type θ. Type θ consumers are endowed with income $z_{\theta t}$ in period t, giving post-transfer income $z_{\theta t}+T_t$. For simplicity, there is no saving or borrowing across periods, so consumers have a period-specific budget constraint $z_{\theta t} + T_t = \boldsymbol{p} \cdot \boldsymbol{q}_t + q_t^n$.

Consumers have quasi-linear flow utility in period t that depends on current consumption and a state variable S_t representing the consumption capital stock from past smoking and vaping. S_t

evolves according to $S_{t+1} = \Lambda(S_t, \boldsymbol{q}_t)$, with Λ increasing in both arguments. Discounted utility from period 0 is

$$U_\theta = \sum_{t=0}^{\infty} \delta^t \left[u_\theta\left(\boldsymbol{q}_t; S_t\right) + q_t^n + z_{\theta t} + T_t\right], \tag{1}$$

where $\delta < 1$ is the discount factor and u_θ is concave in $\boldsymbol{q}_t$. In this general formulation, past consumption S_t can affect both the level of utility (for example, by affecting health) and the marginal utility of consuming c and e (through habit formation). Furthermore, cigarettes and e-cigarettes can be complements or substitutes both in period t and in the long run. For example, they may be substitutes in period t sub-utility u_θ but complements in the long run through effects on S_{t+1}.

Optimizing consumers. Consider first a standard optimizing consumer. Let $V_\theta^*(S_t)$ be the optimizing consumer's value function, after substituting in the period-specific budget constraint. $V_\theta^*(S_t)$ is the solution to the Bellman equation

$$V_\theta^*(S_t) = \max_{\boldsymbol{q}_t} \left[u_\theta(\boldsymbol{q}_t; S_t) - \boldsymbol{p} \cdot \boldsymbol{q}_t + z_{\theta t} + T_t + \delta V_\theta^*(S_{t+1})\right], \tag{2}$$

subject to $S_{t+1} = \Lambda(S_t, \boldsymbol{q}_t)$.

The optimizing consumer's first-order condition for good j is

$$0 = p^j - \left(\frac{\partial u_\theta\left(\boldsymbol{q}_{\theta t}^*; S_t\right)}{\partial q_t^j} + \delta \frac{\partial V_\theta^*\left(S_{t+1}\right)}{\partial S_{t+1}} \cdot \frac{\partial S_{t+1}}{\partial q_t^j}\right), \tag{3}$$

where $\boldsymbol{q}_{\theta t}^*$ denotes optimal consumption for type θ.

Non-optimizing consumers. An important motivation for regulating both cigarettes and e-cigarettes is that consumers may not maximize their utility, perhaps because they have biased beliefs about the health costs of smoking, because they do not correctly predict future habit formation due to forces such as projection bias, or because they are present biased. To model this, we allow consumers to choose $\boldsymbol{q}_{\theta t}$ that differs from $\boldsymbol{q}_{\theta t}^*$ and thus may not maximize utility. These quantities could be derived by assuming that consumers maximize some specific "perceived" utility function such as quasi-hyperbolic utility, but we focus on insights that hold in general for any structural model.[7] Define $V_\theta(S_t) \leq V_\theta^*(S_t)$ as type θ's value function, i.e. the present discounted utility derived from (potentially suboptimal) actual consumption. Substituting in the budget constraint, we can write utility from time t as

$$U_{\theta t}(\boldsymbol{q}_t; S_t) = u_\theta(\boldsymbol{q}_t; S_t) - \boldsymbol{p} \cdot \boldsymbol{q}_t + z_{\theta t} + T_t + \delta V_\theta(S_{t+1}), \tag{4}$$

[7] See Mullainathan, Schwartzstein and Congdon (2012), Chetty (2015), and Bernheim and Taubinsky (2018) for further discussion of the "reduced form" or "sufficient statistic" approach to behavioral public economics.

subject to $S_{t+1} = \Lambda(S_t, \boldsymbol{q}_t)$. Standard optimizing consumers maximize this equation, making it equivalent to Equation (2), but non-optimizing consumers do not.

Following the sin tax literature, we then define bias $\gamma_\theta^j(\boldsymbol{p}, S_t)$ as the difference (in units of dollars) between price and the marginal utility of good j at the chosen consumption levels $\boldsymbol{q}_{\theta t}$:

$$\gamma_\theta^j(\boldsymbol{p}, S_t) := p^j - \left(\frac{\partial u_\theta\left(\boldsymbol{q}_{\theta t}; S_t\right)}{\partial q_t^j} + \delta \frac{\partial V_\theta\left(S_{t+1}\right)}{\partial S_{t+1}} \cdot \frac{\partial S_{t+1}}{\partial q_t^j} \right). \tag{5}$$

Put differently, γ_θ^j is the period t price increase that would induce consumers of type θ to consume $\boldsymbol{q}_{\theta t}^*$. $\gamma_\theta^j > 0$ means that type θ consumes more than the privately optimal amount, $\gamma_\theta^j < 0$ means that type θ consumes less, and $\gamma_\theta^j = 0$ when $\boldsymbol{q}_{\theta t} = \boldsymbol{q}_{\theta t}^*$, per Equation (3). $\gamma_\theta^j(\boldsymbol{p}, S_t)$ depends on prices and consumption in other periods, as these factors affect flow utility and the continuation value function.

To illustrate, consider two examples. First, consider present focused consumers whose smoking and vaping imposes future health harms, in a model with no habit formation. Specifically, assume that $u_\theta(\boldsymbol{q}_t; S_t) = v(\boldsymbol{q}_t) - hS_t$, where the second term is the health harm from past consumption, which evolves according to $S_{t+1} = \rho(S_t + q_t^c + q_t^e)$ for $\rho \in (0,1)$. Considering the infinite discounted sum of future health harms hS_t, the effect of consumption on the continuation value is $\frac{\partial V_\theta(S_{t+1})}{\partial S_{t+1}} \cdot \frac{\partial S_{t+1}}{\partial q_t^j} = -\frac{1}{1-\delta\rho} h \cdot \rho$, so the marginal utility of consumption at $\boldsymbol{q}_{\theta t}$ is $\frac{\partial v(\boldsymbol{q}_{\theta t})}{\partial q_t^j} - \frac{\delta\rho}{1-\delta\rho} h$. Quasi-hyperbolic consumers discount future harms by β_θ, choosing consumption to set $p^j = \frac{\partial v(\boldsymbol{q}_{\theta t})}{\partial q_t^j} - \beta_\theta \frac{\delta\rho}{1-\delta\rho} h$. Substituting marginal utility and the consumption choice into the definition of γ_θ^j from Equation (5) gives

$$\gamma_\theta^j = (1 - \beta_\theta) \frac{\delta\rho}{1 - \delta\rho} h. \tag{6}$$

This is the familiar result that bias is the uninternalized future health cost.[8]

As a second example, imagine that projection bias causes consumers to underestimate habit formation. Specifically, define $\alpha^j := \frac{\partial S_{t+1}}{\partial q_t^j}$ as the habit formation from good j, and allow consumers to misperceive habit formation as $\tilde{\alpha}_\theta^j$. Assume for simplicity that the marginal effect of habit stock on future utility $\frac{\partial V_\theta(S_{t+1})}{\partial S_{t+1}}$ is a constant. The marginal utility of consumption is $\left(\frac{\partial u_\theta(\boldsymbol{q}_{\theta t}; S_t)}{\partial q_t^j} + \delta \frac{\partial V_\theta(S_{t+1})}{\partial S_{t+1}} \cdot \alpha^j \right)$, but consumers choose consumption to set $p_t^j = \left(\frac{\partial u_\theta(\boldsymbol{q}_{\theta t}; S_t)}{\partial q_t^j} + \delta \frac{\partial V_\theta(S_{t+1})}{\partial S_{t+1}} \cdot \tilde{\alpha}_\theta^j \right)$, so

$$\gamma_\theta^j = \delta \frac{\partial V\left(S_{t+1}\right)}{\partial S_{t+1}} \cdot \left(\tilde{\alpha}_\theta^j - \alpha^j \right). \tag{7}$$

Externalities and social welfare. Consumers impose linear externalities $\boldsymbol{\phi}_\theta = \{\phi_\theta^c, \phi_\theta^e\}$ on

[8] Since true utility from Equation (1) uses exponential discounting, this example invokes the long-run criterion, which is not uncontroversial (Bernheim and Rangel 2009; Bernheim and Taubinsky 2018).

the government budget, for example due to increased costs of government-sponsored health care or reduced social security payments due to early death. The results would be the same if some or all of the externality entered other consumers' utility directly, for example due to second-hand smoke. We define $\boldsymbol{\phi}_\theta > 0$ as a negative externality and $\boldsymbol{\phi}_\theta < 0$ as a positive externality. For simplicity, we assume that the externality is imposed in the same period as consumption occurs.

Social welfare from period 0 as a function of taxes $\boldsymbol{\tau}$ is

$$W(\boldsymbol{\tau}) = \sum_\theta s_\theta U_\theta, \tag{8}$$

and the government's balanced budget constraint requires $T_t = \sum_\theta (\boldsymbol{\tau} - \boldsymbol{\phi}_\theta) \cdot \boldsymbol{q}_{\theta t}$ for all t.

1.2 Optimal Taxes

Define the "marginal distortion" φ^j_θ as the sum of the marginal bias and marginal externality for consumer type θ:

$$\varphi^j_\theta(\boldsymbol{p}, S_t) := \gamma^j_\theta(\boldsymbol{p}, S_t) + \phi^j_\theta. \tag{9}$$

$\varphi^j_\theta(\boldsymbol{p}, S_t)$ will be a sufficient statistic for welfare and the optimal tax. This highlights that externalities and internalities enter our model in the same way: they both reflect a difference (in units of dollars) between consumers' perceived marginal utility (revealed by the demand curve) and marginal social welfare.

In Appendix A, we derive socially optimal taxes by maximizing Equation (8) subject to the balanced budget constraint and consumer decision-making. The optimal tax is

$$\tau^{j*} = \underbrace{\frac{\sum\limits_{\theta,t} \delta^t s_\theta \dfrac{dq^j_{\theta t}}{dp^j} \varphi^j_\theta(\boldsymbol{p}, S_t)}{\sum\limits_{\theta,t} \delta^t s_\theta \dfrac{dq^j_{\theta t}}{dp^j}}}_{\text{average marginal distortion}} + \underbrace{\frac{\sum\limits_{\theta,t} \delta^t s_\theta \dfrac{dq^{-j}_{\theta t}}{dp^j} \left(\varphi^{-j}_\theta(\boldsymbol{p}, S_t) - \tau^{-j}_t\right)}{\sum\limits_{\theta,t} \delta^t s_\theta \dfrac{dq^j_{\theta t}}{dp^j}}}_{\text{substitution distortion}}. \tag{10}$$

The first term is the average marginal distortion, familiar from Diamond (1973): the average distortion across types, weighted by each type's own-price response. The optimal tax is larger if the average distortion is larger or if distortions are larger for types who are more responsive to the tax. The second term is a substitution distortion, familiar from Allcott, Lockwood and Taubinsky (2019) and others: the average *uninternalized* distortion from the substitute good, weighted by each type's cross-price response. The optimal tax is larger if a substitute good has a beneficial uninternalized distortion or if a complementary good has a harmful uninternalized distortion.

The demand response $\frac{dq^k_{\theta t}}{dp^j}$ is a total derivative, reflecting changes in period t consumption caused by changes in prices in all periods, including the effects of habit formation. Both $\frac{dq^k_{\theta t}}{dp^j}$ and the marginal distortion $\varphi^j_\theta(\boldsymbol{p}, S_t)$ can vary over time and are affected by changes in tax-inclusive prices and consumption capital stock.

This simple extension of standard formulas has interesting implications in our application. First, the optimal cigarette tax may have changed with the introduction of e-cigarettes. For example, vaping is particularly popular among youth, and youth may have higher marginal internalities and externalities. If there are now fewer youth smokers marginal to the e-cigarette tax, this would decrease the average marginal distortion and thus decrease the optimal cigarette tax. As another example, many states have not yet implemented e-cigarette taxes because vaping is so new. If the average e-cigarette tax is lower than the average marginal distortion and e-cigarettes are substitutes (or complements) for cigarettes, then the substitution distortion from e-cigarettes is negative (positive) and τ^{c*} would decrease (increase). As a final example, e-cigarettes could reduce the health harms from cigarette addiction if addicted cigarette smokers can transition to vaping. With present focus or projection bias, this reduction in the harms from addiction could imply lower bias $\varphi^j_\theta(\boldsymbol{p}, S_t)$ and thus a lower optimal tax τ^{c*}.

A second implication is that the optimal e-cigarette tax could plausibly be negative, i.e. a subsidy, if the substitution distortion from cigarettes is relatively large and negative. This could arise if e-cigarettes are not very harmful (φ^e_θ is small or negative), baseline cigarette taxes are "too low" ($\varphi^c_\theta - \tilde{\tau}^c > 0$), and e-cigarettes are substitutes for cigarettes ($\frac{dq^c_\theta}{dp^e} > 0$).

1.3 Welfare Effect of an E-Cigarette Ban

We model an e-cigarette ban as an increase in the e-cigarette tax from current level $\tilde{\tau}^e$ to ∞ for all periods beginning with period 0. The welfare effect of a ban is thus

$$\Delta W := \int_{\tilde{\tau}^e}^{\infty} \frac{\partial W(\boldsymbol{\tau})}{\partial \tau^e} d\tau^e. \tag{11}$$

If the cigarette and e-cigarette taxes are currently set optimally, then raising τ^e to ∞ by construction reduces welfare in our model. However, a ban may be preferred to taxation for unmodeled reasons such as tax evasion or political constraints on tax rates. We thus allow current taxes $\tilde{\boldsymbol{\tau}}$ to take any value, not necessarily the optimal rates. Furthermore, bias and externalities (and thus optimal tax rates) may vary across types (e.g. youth versus adults), and it may be administratively easier to implement a type-specific ban (e.g. a ban on sales to youth) than to implement type-specific taxes. We thus consider type-specific bans in the welfare analysis in Section 7.

Define $\Delta q^j_{\theta t} := q^j_{\theta t}(\tilde{\tau}^c, \tau^e = \infty) - q^j_{\theta t}(\tilde{\boldsymbol{\tau}})$ as the change in period t consumption of good j from a permanent e-cigarette ban. For e-cigarettes, this is simply period t consumption: $\Delta q^e_{\theta t} = -q^e_{\theta t}(\tilde{\boldsymbol{p}}) < 0$. Further define

$$\overline{\varphi}_{\theta}^{j}(\boldsymbol{p}, S_t) := \frac{\int_{\tilde{\tau}^e}^{\infty} \varphi_{\theta}^{j}(\boldsymbol{p}, S_t) \frac{dq_{\theta t}^{j}}{d\tau^e} d\tau^e}{\Delta q_{\theta t}^{j}}. \tag{12}$$

This is the average distortion over the consumption of good j that is marginal to the e-cigarette ban.

In Appendix A, we show that the welfare effect of a ban is

$$\Delta W = \sum_{\theta,t} \delta^t s_\theta \left[\underbrace{-\int_{\tilde{\tau}^e}^{\infty} q_{\theta t}^{e} d\tau^e}_{\text{perceived CS change}} \underbrace{-\sum_j \Delta q_{\theta t}^{j} \left(\overline{\varphi}_{\theta}^{j}(\boldsymbol{p}, S_t) - \tau^j \right)}_{\text{uninternalized distortion change}} \right]. \tag{13}$$

The first term in Equation (13) is the loss in perceived consumer surplus as traced out by the market demand curve. For standard optimizing consumers, the word "perceived" is unnecessary. We add the word "perceived" to emphasize that with non-optimizing consumers, this term is not the actual change in U_θ that results from the price decrease. The second term captures the change in uninternalized negative distortions from both cigarettes and e-cigarettes. Separating the two terms in this way foreshadows that one can calculate ΔW by estimating perceived consumer surplus with standard demand estimation techniques and then separately quantifying the internalities and externalities in $\overline{\varphi}_{\theta}^{j}$.

The period-specific welfare effects of a permanent ban will change over time as the initial stock of consumption capital $S_{t=0}$ depreciates. For example, reduced S_t in later periods could decrease $q_{\theta t}^{e}$ and make demand more elastic, thereby reducing the perceived consumer surplus loss.

If $\Delta q_{\theta t}^{c} \left(\overline{\varphi}_{\theta t}^{c} - \tau^c \right) = 0$, which holds if e-cigarettes and cigarettes are neither complements nor substitutes or if the cigarette tax exactly internalizes the average distortion marginal to the ban, then the e-cigarette market can be considered in isolation. Otherwise, an e-cigarette ban affects uninternalized distortions in the cigarette market. This effect increases ΔW if $\Delta q_{\theta t}^{c} \left(\overline{\varphi}_{\theta t}^{c} - \tau^c \right) < 0$, which holds if the two products are substitutes ($\Delta q_{\theta t}^{c} > 0$) and the current cigarette tax is "too high" ($\overline{\varphi}_{\theta t}^{c} - \tau^c < 0$) or if the two products are complements and the cigarette tax is "too low." In theory, the reduced uninternalized distortions from cigarettes could justify an e-cigarette ban even if e-cigarettes have no uninternalized distortions. This is analogous to arguments for banning drugs like marijuana on the grounds that they are not particularly harmful on their own but could be gateways to more harmful drugs.

1.4 Empirical Implementation

For empirical implementation, we define a substitution parameter $\sigma_{\theta t} := \frac{dq_{\theta t}^{c}/dp^e}{dq_{\theta t}^{e}/dp^e}$ representing the ratio of demand responses to a permanent e-cigarette price change. We further define $\varphi_{\theta}^{j} =$

$\mathbb{E}_t\left[\varphi^j_\theta(\boldsymbol{p}, S_t)|\theta\right]$, $\sigma_\theta := \mathbb{E}_t\left[\sigma_{\theta t}|\theta\right]$, and $q^j_\theta := \mathbb{E}_t\left[q^j_{\theta t}|\theta\right]$ as expectations over time. σ_θ captures the net long-run substitutability between e-cigarettes and cigarettes.

We also impose two assumptions. First, we assume a homogeneous and time-invariant own-price elasticity $\eta := \frac{\partial q^e_{\theta t}}{\partial p^e}\frac{p^e}{q^e_{\theta t}} < 0$, because the Nielsen RMS data do not allow us to separately estimate elasticities by consumer type. Second, we assume pairwise zero covariance between the marginal distortion $\varphi^j_\theta(\boldsymbol{p}, S_t)$, substitution $\sigma_{\theta t}$, consumption $q^j_{\theta t}$, and time t for each type. While this assumes away potentially interesting dynamics, we are not able to credibly estimate how any of these parameters covary or would change over time in response to a tax or ban.

Under these assumptions, Appendix A shows that the optimal tax formula from Equation (10) reduces to

$$\tau^{e*} = \frac{\sum_\theta s_\theta q^e_\theta \left[\varphi^e_\theta + \sigma_\theta\left(\varphi^c_\theta - \tilde{\tau}^c\right)\right]}{\sum_\theta s_\theta q^e_\theta}, \tag{14}$$

where the first term inside the brackets is the e-cigarette marginal distortion, and the second term is the uninternalized substitution distortion from cigarettes.

To estimate the welfare effect of a ban, we write the expected cigarette consumption change as $\Delta q^c_\theta = -\sigma_\theta q^e_\theta(\tilde{\boldsymbol{p}})$. To estimate perceived consumer surplus change, some assumption is required because observed market prices do not rise high enough to identify the demand function at high prices. This identification problem and the use of functional form assumptions such as linear or logit demand are common in related literature (Hausman 1996; Petrin 2002). We assume that each type's perceived consumer surplus change equals the area under a linear demand curve drawn tangent to their demand function at current prices, which is the triangle $\Delta q^e_\theta \frac{\tilde{p}^e}{-2\eta} < 0$.

Under these assumptions, Appendix A shows that the welfare effect of an e-cigarette ban in the average period is

$$\Delta\overline{W} = \sum_\theta s_\theta \left[\underbrace{\Delta q^e_\theta \frac{\tilde{p}^e}{-2\eta}}_{\text{perceived CS change}} - \underbrace{\sum_j \Delta q^j_\theta\left(\varphi^j_\theta - \tau^j\right)}_{\text{uninternalized distortion change}}\right]. \tag{15}$$

In the rest of the paper, we estimate τ^{e*} and $\Delta\overline{W}$ using these formulas.

2 Data

2.1 Nielsen Scanner Data

For our price elasticity estimates in Section 4, we use scanner data from Nielsen's Retail Measurement Services (RMS). The data include weekly prices and sales volumes by UPC at approximately 27,000 stores in the contiguous U.S. from 96 retail chains. RMS includes e-cigarette products beginning in 2013, and 2017 is the most recent year currently available. See Appendix B for RMS data construction details.

RMS includes 53, 32, 55, and 2 percent of total sales in the grocery, mass merchandiser, drug, and convenience store channels, respectively. In addition to its very limited coverage of convenience stores, RMS has no coverage of vape shops or online channels where many e-cigarette products are sold. In 2017, RMS stores sold \$114 million in e-cigarette products, out of the \$4.6 billion sold nationwide as shown in Figure 1. This 2.5 percent coverage rate is an important limitation of the data.[9]

We collected data on the volume of each UPC (in milliliters of e-liquid) from online databases, manufacturer websites, store visits, and from a database kindly shared by the authors of Cotti et al. (2020).

As shown in Appendix Table A1, 11 states, counties, or cities in the contiguous U.S. initiated or changed e-cigarette taxes between 2013 and 2017. We use these tax changes for identification. For our empirical analysis, we define 51 geographic "clusters": the two counties (Montgomery County, Maryland and Cook County, Illinois) that have county-level e-cigarette taxes, the contiguous 48 states (where Maryland and Illinois exclude Montgomery County and Cook County), and Washington, D.C.[10] We collapse the UPC-store-week RMS data to the level of UPC-cluster-month, calculating total units sold and quantity-weighted average price.

2.2 Smoking and Vaping Sample Surveys

For our substitution estimates in Section 5, we use all major annual surveys that have recorded information on vaping and/or smoking for adults and/or youth in the U.S. since 2004: the Behavioral Risk Factor Surveillance System (BRFSS), the National Health Interview Survey (NHIS), the National Survey of Drug Use and Health (NSDUH), Monitoring the Future (MTF), and the National Youth Tobacco Survey (NYTS). Table 1 presents information on each dataset. We have 7.4 million observations across the five datasets in total, or about 500,000 per year. All estimates

[9]Although the household-level Nielsen Homescan data could also be useful in exploring heterogeneity and measuring additional purchases outside of RMS stores, Homescan's effective sample size is much smaller: Homescan, with 60,000 households, covers about 0.05 percent of the U.S., against the 2.5 percent in RMS.

[10]The city of Chicago also has an e-cigarette tax; we add this to the Cook County tax because the RMS store data include identifiers for county but not city.

in the paper are weighted for national representativeness for adults (people aged 18 or older) and youth (people in grades 6-12).

Appendix B details how we construct consistent smoking and vaping variables. We construct smoking in units of packs of cigarettes smoked per day and vaping in units of share of days vaped. In all datasets other than BRFSS, we can directly estimate the number of packs per day smoked. BRFSS only records whether someone smokes or vapes "every day," "some days," or "not at all," but we use conditional means from the other adult datasets to impute packs per day smoked and share of days vaped. The datasets do not include the quantity of e-liquid used or the nicotine content of cigarettes or e-liquid.

Demographic variables are central to our analysis. From the possible set of standard demographics (age, race/ethnicity, etc.), we include a demographic variable only if it is observed consistently across all datasets. We denote the vector of demographic group indicators for person i as $\boldsymbol{G}_i$. For adults, $\boldsymbol{G}_i$ includes race/ethnicity (Asian, Black, other/missing, Hispanic, White), sex (male/female), educational attainment (high school, less than high school, some college, college graduate), income quintiles, and age groups (18–24, 25–29, 30–49, 50–64, and 65+). For youth, $\boldsymbol{G}_i$ includes race (Black, other/missing, Hispanic, White), sex, and each grade from 6–12.[11] We refer to demographic "cells" as the interactions of our demographic group indicators, e.g. "Asian women aged 18–24 who are college graduates and are in the lowest-income quintile." There are $5 \times 2 \times 4 \times 5 \times 5 = 1{,}000$ cells for adults and $4 \times 2 \times 7 = 56$ cells for youth.

In our regressions, we will include "dataset controls" to address two sampling issues. First, in 2011, BRFSS was updated to sample people using cell phones instead of only people with land lines (Pierannunzi et al. 2012). This causes an artificial change in smoking rates, and this change could differ across demographic groups. Second, the NYTS is collected in 2004, 2006, 2009, and annually since 2011, but not in 2005, 2007, 2008, or 2010.

2.3 E-cigarette User Survey

To estimate the average e-liquid price and quantity consumed per day, we ran a survey we call the E-cigarette User Survey in August 2019. The sample is an online panel of U.S. e-cigarette users provided by polling firm SurveyMonkey through their Audience Panel service. We asked whether people now use e-cigarettes every day, some days, or not at all, the number of days vaped out of the past 30, the milliliters of e-liquid consumed in the past 30 days, and the amount of money they spent to buy the e-liquid consumed in the past 30 days.[12] We have 147 valid responses. We weight the sample to be representative of U.S. adults who vaped in the past 30 days on income, gender, and vaping frequency.

[11] We are limited to four race/ethnicity groups in the youth dataset because Asian is not a separate category from other race in the public-use MTF.

[12] The survey instrument can be accessed from https://www.surveymonkey.com/r/YRZSZZY.

We estimate that the average e-liquid price is $\tilde{p}^e \approx \$3.90$ per milliliter (ml). For comparison, the popular 0.7 milliliter Juul pods cost \$6.41/ml at average tax rates, while large 100 ml e-liquid bottles can be as cheap as \$0.50/ml. The average day of vaping involves $\Gamma \approx 0.58$ milliliters of e-liquid consumption, slightly less than one Juul pod. This is more than the unweighted average across vapers of consumption per day, because people who vape every day consume more e-liquid per day than people who vape on some days.

2.4 E-cigarette Expert Survey

To quantify the state of knowledge on the health effects of vaping, we carried out a survey we call the E-cigarette Expert Survey in late January 2020. The sampling frame was the 123 corresponding authors of papers on the health impacts of e-cigarettes that are cited in the landmark National Academy of Sciences (2018) report (excluding those with corporate affiliations) and the 43 authors of papers about cigarettes or e-cigarettes cited in Cutler et al. (2015), Chaloupka, Levy and White (2019), and our September 2019 draft. We received 55 responses, 38 from the National Academy sample and 17 from the economist sample, giving a response rate of 33 percent.[13]

The survey asked experts to consider "average adults who would say they vape 'every day' compared to average adults who would say they smoke cigarettes 'every day,' consuming comparable amounts of nicotine." We instructed the experts to "not include vaping or smoking of THC/marijuana." After explaining and giving examples of externalities and internalities, we asked five questions:

1. Imagine that every-day cigarette smoking by the average adult imposes 100 units of external harms. In comparison, how large do you think are the external harms from every-day vaping?

2. Now imagine that every-day cigarette smoking by the average adult imposes 100 units of misperceived internal harms. In comparison, how large do you think are the misperceived internal harms from every-day vaping?

3. Now imagine that every-day vaping imposes 100 units of external and internal harms for average adults. In comparison, how large do you think are the harms from every-day vaping for average youth?

4. Do you think banning all e-cigarettes is a good idea?

5. Do you think banning flavored e-cigarette products for all users is a good idea?

Both negative and positive answers were allowed for the first three questions, although all responses we received were strictly positive.

[13] The survey instrument can be accessed from https://nyu.qualtrics.com/jfe/form/SV_3J0zPeyGgeuU0Hr. We sent the survey to the National Academy and economist samples on different days, allowing us to infer with high probability which sample a survey response came from.

3 Smoking and Vaping Trends

Figure 1 presents trends in U.S. sales of cigarettes and e-cigarettes. Cigarette sales decreased by 40 percent (from 20 billion to 12 billion packs) from 2004 to 2018. While the first modern e-cigarettes became available in the late 2000s, sales were relatively low until about 2013. Sales grew continually from 2013 to 2017 and increased notably in 2018 with the introduction of the popular Juul e-cigarette.

Figure 2 presents trends in smoking and vaping recorded in the sample surveys. Self-reported adult smoking in Panel (a) declined by about 45 percent (from about 0.15 to 0.08 packs per adult per day) from 2004 to 2018. The 2011 jump in the BRFSS trend is due to the sampling frame change discussed earlier. Youth smoking in Panel (b) dropped by an even larger proportion, from about 0.035 to less than 0.01 packs per youth per day.

Prior work has established that the levels and trends line up imperfectly between these two figures. In Appendix B.2.8, we calculate that the sample survey data overstate e-cigarette sales and understate cigarette sales by an amount consistent with earlier estimates by Liber and Warner (2018). The 2004–2018 percent smoking reductions are fairly consistent between the two figures. Self-reported vaping grew much less quickly than the e-cigarette sales data, although the 2018 increase in self-reported youth vaping is consistent with the 2018 sales increase.

On the cigarette consumption figures, we add a vertical line to mark the time just before e-cigarette sales started to take off in 2013. The smoking declines in Figures 1 and 2 are close to linear, with no substantial changes as e-cigarettes became popular after 2013. Unless there was some countervailing force that would have changed cigarette consumption at the same time that vaping became popular, this suggests that e-cigarettes had little impact on overall cigarette consumption. Levy et al. (2019) make a similar point focusing on youth vaping.

To quantify this idea, recall the substitution parameter $\sigma_\theta = \mathbb{E}_t \left[dq^c_{\theta t}/dq^e_{\theta t}|\theta\right]$, in units of cigarette packs per day vaped. The introduction of e-cigarettes increases $q^e_{\theta t}$ from 0 to $q^e_{\theta t}(\tilde{\boldsymbol{p}})$, which in turn changes cigarette consumption by $\sigma_\theta q^e_{\theta t}(\tilde{\boldsymbol{p}})$. In the sample survey data, the average day of smoking by an adult (youth) involves 0.5 (0.15) packs smoked. Thus, $\sigma_\theta \approx -0.5$ ($\sigma_\theta \approx -0.15$) implies that the average smoking day and the average vaping day are perfect substitutes for adults (youth), and $\sigma_\theta \approx 0.5$ ($\sigma_\theta \approx 0.15$) implies that they are perfect complements for adults (youth).

An average vaping day costs 0.58ml $\times$ \$3.90/ml $\approx$ \$2.26 of e-liquid, so if the \$6.9 billion in 2018 e-cigarette sales were all for e-liquid, this would be equivalent to 3.05 billion average vaping days. At 0.5 cigarette packs per average smoking day, 3.05 billion average smoking days would equal about 1.5 billion packs. Thus, if the average vaping and average smoking days were perfect complements (substitutes) over a several-year horizon, cigarette sales would have increased (decreased) by 1.5 billion packs per year by 2018 relative to a counterfactual without e-cigarettes. Since the sales decline on Figure 1 is close to linear over 2004–2018, daily vaping and daily smoking could therefore only be perfect complements or perfect substitutes if the counterfactual sales trend would have been noticeably different from its long-standing historical pattern.

We can do a similar exercise for the sample survey data in Figure 2. In each panel, the left and right y-axes have the same scales. Panel (a) shows that adults vaped on share 0.025 of days in 2018. Thus, if $\sigma_\theta = 0.5$ (or $\sigma_\theta = -0.5$) over several years, adult cigarette consumption would have increased (or decreased) by about 0.0125 packs per day relative to counterfactual. Since the adult cigarette consumption decline on Panel (a) is close to linear over 2004–2018, σ must be relatively close to zero unless the counterfactual smoking trend would have changed noticeably after 2013. This visual argument is particularly clear for youth, who vape on share 0.05 to 0.08 of days in 2018 but have a steady linear decline in cigarette consumption to less than 0.01 packs per day by 2018.

Of course, this visual argument relies on strong assumptions about counterfactual trends and cannot easily rule out values of σ closer to 0. We build on this intuition for a more precise estimate of σ in Section 5.

4 Price Elasticity

4.1 Empirical Strategy

In this section, we use tax changes to estimate the own price elasticity η and the substitution parameter σ_θ using Nielsen RMS data. We index UPCs by k, geographic clusters by s, and months by t. Let q^e_{kst}, $\tilde{p}^e_{kst}$, and $\tilde{\tau}^e_{kst}$ denote quantity sold, sales-weighted average tax-inclusive price, and the ad-valorem tax rate, respectively, for e-cigarette UPCs. Let $\tilde{p}^c_{st}$ and $\tilde{\tau}^c_{st}$ denote the sales-weighted average tax inclusive price and average tax rate as a percentage of tax-exclusive price, respectively, for cigarettes in a given state and month.[14] Let $\boldsymbol{X}_{st}$ denote a cluster-specific linear time trend and an additional vector of controls for potential confounders that might be correlated with both taxes and consumption: the state unemployment rate and beer tax rate as well as indicators for whether the state has an indoor vaping ban, has a medical marijuana law, has passed or implemented a prescription drug program, and implemented the Medicaid expansion.

Let E_{0st} be an indicator variable that takes value 1 if month t is 0–2 months after an e-cigarette tax change in cluster s, and define the vector $\boldsymbol{Q}_{kst} = [E_{0st}, E_{0st}\ln(\tau^e_{kst}+1)]$. The event study figure presented below suggest that prices and sales are slow to adjust in the first quarter after a tax change; controlling for $\boldsymbol{Q}_{kst}$ identifies the elasticity η beginning in the second quarter. Finally, let ν_{kt}, μ_{ks}, and ξ_{st}, respectively denote UPC-month, UPC-cluster, and census division-month fixed effects.

Our primary specification is

[14] Some e-cigarette taxes are "specific" taxes per milliliter of e-liquid, and all cigarette taxes are specific taxes per pack. We transform these tax rates to the implied ad-valorem rate using the UPC's size and price. See Appendix B for details.

$$\ln(q^e_{kst}+1) = \eta \ln(\tilde{p}^e_{kst}) + \chi^e \ln(\tilde{p}^c_{st}) + \boldsymbol{\beta X}_{st} + \boldsymbol{\kappa Q}_{kst} + \nu_{kt} + \mu_{ks} + \xi_{st} + \varepsilon_{kst}, \tag{16}$$

where we instrument for $\ln(\tilde{p}^e_{kst})$ and $\ln(\tilde{p}^c_{st})$ with $\ln(\tilde{\tau}^e_{kst}+1)$ and $\ln(\tilde{\tau}^c_{st}+1)$. The coefficient η is our estimate of the own-price elasticity of demand for e-cigarettes. The coefficient χ^e is the elasticity of substitution, which we transform into σ_θ below.

We keep the estimates at the UPC level instead of aggregating for two reasons. First, unlike cigarettes, there is no individual unit that is natural to aggregate across UPCs: vapor products are primarily e-liquid refills but also include e-cigarette base units and starter kits with both base units and e-liquid. Second, the price variation across UPCs provides additional variation to identify the effects of ad-valorem taxes. To estimate the aggregate elasticity of demand for e-cigarettes using UPC-level data while accounting for the possibility that elasticities might vary by UPC, we would like to weight observations by sales. To avoid mechanical biases arising from the effect of taxes on sales, we weight each UPC-cluster-month observation by the UPC's sales in non-taxed clusters in that calendar year, normalized by total sales across all UPCs in non-taxed clusters in that year. We cluster standard errors by geographic cluster.

We also present event study figures to test for any trends before tax changes and examine how the tax effects vary over time. In four geographic clusters, e-cigarette tax rates change twice during the sample period. We index tax change events within a cluster by $v \in \{1,2\}$, and we define $\mathcal{V}_s$ as the set of changes within cluster s. We define $\Delta \ln(\tilde{\tau}_{ksv}+1)$ as the change in the log e-cigarette tax variable that occurs for UPC k in cluster s in event v. Let E_{qst} represent an indicator variable that takes value 1 if month t is q quarters after an e-cigarette tax change in cluster s, with E_{0st} as defined above.[15] We then estimate

$$y_{kst} = \sum_{v\in\mathcal{V}_s}\sum_{q\in\mathcal{Q}} \eta_q E_{qst} \Delta \ln(\tilde{\tau}_{ksv}+1) + \chi^e \ln(\tilde{\tau}^c_{st}+1) + \boldsymbol{\beta X}_{st} + \nu_{kt} + \mu_{ks} + \xi_{st} + \varepsilon_{kst}, \tag{17}$$

for $y_{kst} \in \{\ln(q_{kst}+1), \ln \tilde{p}_{kst}\}$. Since we have μ_{ks} fixed effects and $\Delta \ln(\tilde{\tau}_{ksv}+1)$ is constant within ks for each tax change event, we let $\mathcal{Q}$ be a mutually exclusive and exhaustive set of event time indicators excluding -1 (the quarter before the tax change) to avoid collinearity.

This empirical strategy has several limitations. First, as we have discussed, RMS covers only 2.5 percent of national e-cigarette sales. The demand elasticity estimated in RMS might differ from the true nationwide demand elasticity if RMS stores serve a non-representative set of e-cigarette consumers or if consumers substitute toward or away from RMS stores in response to a tax. For example, consumers might substitute purchases to retailers in other states or to illegal retailers that evade taxes. Second, while we observe sales for up to several years after a tax change, our estimates

[15] Specifically, $E_{1st} = 1$ if month t is 3–5 months after a tax change, $E_{2st} = -1$ if month t is 1–3 months before a tax change, etc.

may still not reflect the full long-run price elasticity if habit formation takes longer to manifest. Third, we must assume that no other factors affected e-cigarette demand at the same time as the tax changes. Rees-Jones and Rozema (2020) show that local media coverage of cigarettes increases as cigarette taxes are debated and implemented, and such forces could also change e-cigarette demand as e-cigarette taxes are implemented.

4.2 Event Study Figures

Panels (a) and (b) of Figure 3 presents estimates of Equation (17) with $\ln \tilde{p}^e_{kst}$ and $\ln(q^e_{kst}+1)$ as the dependent variables. Panel (a) shows that we have a strong first stage: in the six quarters after a tax change, retail prices rise by 0.5–0.8 log points, suggesting substantial but not full pass-through of the tax. Panel (b) presents the reduced form: in the six quarters after a tax change, quantities decline by 0.7–1.2 log points. We can divide these first stage and reduced form coefficients for an approximate IV estimate of $\eta \approx -1.5$, although this approximation to a Wald estimator only holds if the other endogenous variable (cigarette price) has no effect on e-cigarette price or demand. There is no trend in either prices or quantities in the six quarters before the tax change. Appendix Figure A2 shows that we get very similar point estimates and more precise standard errors when we exclude the cluster-specific linear time trends.

4.3 Parameter Estimates

Table 2 presents estimates of Equation (16). Panel (a) presents the first stages and reduced form. Columns 1 and 2 show that a tax on one good strongly predicts that good's price, but not the other good's price. Column 3 shows that e-cigarette taxes reduce e-cigarette demand, while cigarette taxes have a positive but insignificant coefficient.

Panel (b) presents the instrumental variables estimates of η and χ^e. Our primary estimate in column 1 suggests that e-cigarette demand is more than unit elastic, with $\hat{\eta} \approx -1.32$. Columns 2–4 progressively add fixed effects; after the UPC-cluster effects, the additional fixed effects make little difference. Column 5 presents the primary estimates without the cluster-specific linear time trend; this reduces the estimate to $\hat{\eta} \approx -1.00$. Column 6 shows that the additional controls in $\boldsymbol{X}_{st}$ make little difference. Column 7 presents estimates in a "quasi-panel" in which we add zero-sales observations for all UPCs that had non-zero sales in cluster s in any prior month, but the panel begins with the first month in which we observe any sales in that UPC-cluster. We impute price $\tilde{p}^e_{kst}$ from the last month a sale was observed in that cluster. This also does not substantially change the estimates.

In column 1, the point estimate of the substitution elasticity is $\hat{\chi}^e \approx 0.27$, with standard error of 0.46. In the other columns, $\hat{\chi}^e$ is more positive but not statistically different at conventional levels. Column 5 shows that excluding the cluster-specific linear time trends gives $\hat{\chi}^e \approx 0.66$, with a standard error of 0.34. Appendix Table A2 presents symmetric estimates of cigarette demand on cigarette and e-cigarette prices (instrumented by taxes), using an equation analogous to Equation (16). The resulting substitution parameter is $\chi^c \approx -0.06$, with a standard error of 0.30. Excluding the cluster-specific linear time controls gives $\hat{\chi}^c \approx 0.77$, with a standard error of 0.26. Appendix Figure A3 shows that without these linear time controls, there is an upward trend in cigarette purchases in the six quarters before the e-cigarette tax change. If that upward trend would have continued after the tax change, this would produce an upward-biased estimate of the cross-price elasticity χ^c. This is why we include the cluster-specific linear time controls in our primary specification.[16]

Appendix Tables A3 and A4 present additional robustness checks. The price elasticity estimates do not change substantially if we limit the identification of η to the 18-month window around the tax change, exclude e-cigarette UPCs with imputed volumes, or include only clusters with ad-valorem taxes, excluding clusters with specific taxes. When we exclude the controls $\boldsymbol{Q}_{kst}$ and thereby also identify off of the effects in the quarter beginning with the tax change, the e-cigarette $\hat{\eta}$ estimate moves slightly toward zero. This is consistent with the small quantity effect in quarter $q = 0$ shown in Panel (b) of Figure 3.

We can use the cross-price elasticities to estimate the average substitution parameter σ. Beginning with χ^e from Table 2 and using Slutsky symmetry and quasi-linear demand, we have a population average substitution parameter

$$\sigma = \frac{\partial q^c_\theta / \partial p^e}{\partial q^e_\theta / \partial p^e} = \frac{\partial q^e_\theta / \partial p^c}{\partial q^e_\theta / \partial p^e} = \frac{\chi^e}{\eta} \frac{\tilde{p}^e \Gamma}{\tilde{p}^c}, \tag{18}$$

where Γ converts $\tilde{p}^e$ to units of dollars per day vaped, giving σ in the desired units of packs of cigarettes per day vaped. This gives $\hat{\sigma} \approx -0.059$ (standard error ≈ 0.091), consistent with mild substitutability. Similarly, beginning with χ^c from Appendix Table A2, we have

$$\sigma_\theta = \frac{\partial q^c_\theta / \partial p^e}{\partial q^e_\theta / \partial p^e} = \frac{\chi^c}{\eta} \frac{q^c_\theta}{q^e_\theta}. \tag{19}$$

Using q^c_θ and q^e_θ from the sample survey data displayed in Figure 2, this gives $\hat{\sigma}_{\text{youth}} \approx 0.005$ (SE ≈ 0.027) and $\hat{\sigma}_{\text{adult}} \approx 0.144$ (SE ≈ 0.766). Combining these two estimates using a minimum distance estimator gives $\sigma_{\text{youth}} \approx -0.000$ (almost exactly zero, with SE ≈ 0.026) and $\hat{\sigma}_{\text{adult}} \approx -0.056$ (SE ≈ 0.090). See Appendix D.1 for additional details.

These substitution parameter estimates are credible because they are identified from plausibly exogenous tax changes in administrative data. However, we have seen that the point estimates are somewhat sensitive to controls, and we are not able to estimate separate substitution elasticities for

[16] Using analogous regressions in the RMS data, Cotti et al. (2020) estimate an e-cigarette own-price elasticity of -1.5, closely in line with our estimates. They do not include linear time trends in any specification, and their cross-price elasticity estimates χ^e and χ^c are approximately 1, indistinguishable from our estimates from specifications excluding cluster-specific linear time trends.

youth versus adults. An additional alternative approach to estimating the substitution parameter σ_θ would therefore be valuable.

5 Substitution Patterns

5.1 Graphical Illustrations

In this section, we extend the graphical analysis of cigarette smoking trends from Section 3 into a formal empirical strategy for estimating the substitution parameter σ. While Section 3 considered aggregate nationwide data, we now exploit the fact that e-cigarette demand varies substantially across demographic groups.

To demonstrate this demand variation, we regress e-cigarette use on a vector demographic group indicators $\boldsymbol{G}_i$ using the following equation:

$$q_{it}^e = \boldsymbol{\kappa}\boldsymbol{G}_i + \xi_{it}^e, \tag{20}$$

where i indexes individuals in the sample surveys and t indexes years. Figure 4 presents results for adults and youth. White people (the omitted race category), men, non-college graduates, lower-income people, and younger adults (but older youth) have higher e-cigarette demand.[17]

What explains this variation? Academic papers (Hartwell et al. 2017; Pepper et al. 2014; Perikleous et al. 2018) and industry sources (Bour 2019) discuss early adopters of e-cigarettes and often draw analogies to early adopters of other technologies. To explore this, Appendix Figure A5 presents estimates of Equation (20) for social media use in 2008 and internet use in 2000. As with e-cigarettes, men and younger adults were more likely to adopt these other new technologies. One difference is that people with less formal education are conditionally more likely to vape, whereas they were conditionally less likely to be early adopters of social media and the internet.

Figure 5 presents smoking and vaping trends for people with above- versus below-median predicted vaping $\hat{\boldsymbol{\kappa}}\boldsymbol{G}_i$. Cigarette use is residual of dataset controls that address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. As in Figure 2, the right and left panels are on the same scales. The figures show that high-vaping demographics are also high-smoking demographics, and the high-smoking demographics are reducing smoking faster than low-smoking demographics. For both high- and low-vaping demographics, smoking is decreasing at a very steady annual rate beginning in 2004.

The vertical red line before 2013 again marks the time when e-cigarette sales start to take off. Recall that in the sample survey data, the average day of smoking by an adult (youth) involves 0.5 (0.15) packs smoked. If an average day of vaping were a perfect complement (or perfect substitute)

[17] Appendix Figure A4 shows that these patterns are similar across the multiple datasets that record vaping, although the estimated coefficients vary slightly.

for an average day of smoking, one would expect that the relative cigarette consumption of high-vaping demographics would start to increase (or decrease) after 2013. In reality, is difficult to visually detect any change in the annual smoking decreases as e-cigarettes become popular.

Figure 6 continues this logic by presenting the difference in cigarette use between the same high- and low-vaping demographics. The dashed line is a time trend fitted only on pre-2013 data, while the solid line is a time trend fitted only on post-2013 data. The top (bottom) of the shaded area at the right of the figure presents the predicted difference in smoking if $\sigma_\theta = 1$ ($\sigma_\theta = -1$), i.e. if daily vaping were a perfect complement (perfect substitute) for smoking one pack per day.[18] For adults, the actual smoking difference is slightly below the pre-2013 prediction until 2018, but much closer to zero than to the $\sigma_\theta = -1$ bound. This suggests limited complementarity or substitutability. For youth, the actual smoking difference is almost exactly the same as the pre-2013 prediction, suggesting close to zero complementarity or substitutability.

Appendix Figures A6–A9 present versions of Figure 5 for splits of each specific demographic group (sex, race, age/grade, education, and income). Appendix Figures A10–A13 present versions of Figure 6 for the most predictive split of each demographic group (e.g. Whites versus non-Whites, college versus non-college adults, etc.). These allow informal overidentification tests. The results are quite similar across all groups, suggesting limited complementarity or substitutability, with one exception: the adult income split suggests very strong complementarity, as there is little difference in vaping by income, and thus even small deviations from trends by income group are large when scaled by the small vaping difference. In econometric terms, this means that income is a relatively weak instrument. This will not drive the formal estimates below because other vaping predictors generate more variation in $\hat{q}^e_{it}$.

5.2 Empirical Strategy

To identify σ_θ, we estimate the extent to which the demographic differences in vaping affect cigarette consumption after e-cigarettes are introduced. Specifically, we regress cigarette consumption on e-cigarette consumption using two-stage least squares (2SLS), instrumenting for e-cigarette consumption with demographic-by-time predictors and controlling for linear time trends. Let ν_t denote year indicators, and let μ_{dgt} denote "dataset controls" to address the 2011 BRFSS sampling change and the fact that NYTS is not available in certain years.[19] The second stage regression is

[18]To construct the perfect complement (substitute) predictions, we predict smoking using the pre-2013 time trend and then add (subtract) average vaping in the years when it is observed. Specifically, define $\hat{q}^c_{Ht}$ and $\hat{q}^c_{Lt}$ as the predicted smoking rates for people in high- and low-vaping demographics, and define q^e_{Ht} and q^e_{Lt} as their actual vaping rates in year t. The perfect complement and substitute bounds for group $g \in \{H, L\}$ are $\hat{q}^c_{gt} \pm q^e_{gt}$. The bounds plotted on the figure are $(\hat{q}^c_{Ht} - \hat{q}^c_{Lt}) \pm (q^e_{Ht} - q^e_{Lt})$.

[19]For adults and youth, $\boldsymbol{\nu}_{dgt}$ includes an indicator for each dataset (with NSDUH as the omitted dataset) interacted with the demographic indicators $\boldsymbol{G}_i$. For adults, $\boldsymbol{\nu}_{dgt}$ also includes a pre-2011 indicator and a pre-2011 BRFSS indicator, both interacted with $\boldsymbol{G}_i$. The $\boldsymbol{\nu}_{dgt}$ controls thereby address the variability introduced by BRFSS and NYTS sampling and rescale smoking to levels in the NSDUH.

$$q^c_{it} = \sigma \hat{q}^e_{it} + \boldsymbol{\lambda} \boldsymbol{G}_i + \boldsymbol{\omega}(t-2004)\boldsymbol{G}_i + \nu_t + \mu_{dgt} + \varepsilon_{it}. \tag{21}$$

The inclusion of group-specific intercepts and time trends $\boldsymbol{G}_i$ and $(t-2004)\boldsymbol{G}_i$ mean that we identify σ_θ from changes in smoking conditional on those linear trends. However, because we now exploit demand variation across demographic groups, we can also include time dummies ν_t that soak up demand shifts that are common across groups in levels, although not in proportions.

The instruments for vaping q^e_{it}, denoted $\boldsymbol{Z}_{it}$, are $\boldsymbol{G}_i \cdot 1[t \geq 2013]$, $\boldsymbol{G}_i \cdot 1[t \geq 2013] \cdot (t-2012)$, and $\boldsymbol{G}_i \cdot 1[t = 2018]$, where $1[\cdot]$ denotes the indicator function. The first two sets of instruments allow vaping to have different levels and trends by demographic group after vaping begins to grow in 2013. The third set is useful in fitting the 2018 increase in youth vaping seen in Figure 2.

The first stage is

$$\tilde{q}^e_{it} = \boldsymbol{\zeta} \mathbf{Z}_{it} + \boldsymbol{\lambda}^1 \boldsymbol{G}_i + \boldsymbol{\omega}^1(t-2004)\boldsymbol{G}_i + \nu^1_t + \mu^1_{dgt} + \varepsilon_{it}, \tag{22}$$

where e-cigarette consumption $\tilde{q}^e_{it}$ is defined below, and we use "1" superscripts to indicate first-stage parameters.

We must modify the first stage for two reasons. First, q^e_{it} is not recorded in any dataset for the years between when e-cigarettes were introduced and 2014 (for youth) and 2016 (for adults). We denote this initial year with vaping data as $\underline{t}$. Second, q^e_{it} is not recorded at all in the NSDUH data, and it is missing for about ten percent of observations in dataset-years when it is supposed to be recorded.

To address the missing q^e_{it} for early years, we impute the averages by demographic group assuming linear growth from zero in 2012 to the level in year $\underline{t}$. This assumption is motivated by the sales trends from Figure 1, which showed limited vaping until 2013 and roughly linear growth for the several years after that. We predict vaping by demographic group by estimating Equation (20) with data from year $\underline{t}$, giving demographic coefficients $\hat{\boldsymbol{\kappa}}_{\underline{t}}$, and then construct observed or imputed vaping as follows:

$$\tilde{q}^e_{it} = \left\{ \begin{array}{cc} q^e_{it}, & t \geq \underline{t} \\ \hat{\boldsymbol{\kappa}}_{\underline{t}} \boldsymbol{G}_i \cdot \frac{t-2012}{\underline{t}-2012}, & 2013 \leq t < \underline{t}-1 \\ 0, & t \leq 2012 \end{array} \right\}. \tag{23}$$

We carry out this imputation in all datasets other than NSDUH.

To address the missing vaping data in the NSDUH (for all years) and in other datasets (beginning in year $\underline{t}$), we use two-sample 2SLS. We estimate the first stage (Equation (22)) in all datasets other than NSDUH, construct the fitted values $\hat{q}^e_{it}$ for all observations, and run the second stage (Equation (21)) with all observations.[20] We bootstrap the entire procedure including imputation steps and

[20]We impute predicted values with dataset controls for the NSDUH by assuming that NSDUH is the average of

draw bootstrap samples by demographic cell.

This approach is a cousin of the "shift-share" identification strategy popularized by Bartik (1991) and Blanchard and Katz (1992), and discussed in Goldsmith-Pinkham, Sorkin and Swift (2019): we primarily exploit cross-sectional variation in demand across demographic groups with the time-series growth of e-cigarette use. The exclusion restriction is that the instruments affected post-2013 smoking only through vaping—intuitively, that there would have been no changes in smoking trends for higher- versus lower-vaping demographics if e-cigarettes had not been introduced.

We provide two types of suggestive evidence in favor of the exclusion restriction. First, we present a set of informal overidentification tests using different demographic groups as instruments. If the estimates remain stable across different demographic groups, then any potential confounder must have affected all demographic groups. Second, we present graphical event studies that test for trends in smoking in demographics with high versus low latent e-cigarette demand, *before* e-cigarettes were introduced. If there are no such trends, then any potential confounder must have arisen at the same time as e-cigarettes became popular.

The event study regression is analogous to our second stage (Equation (21)), except that ζ is allowed to vary by year:

$$q_{it}^{c} = \boldsymbol{\zeta}_t \left(\hat{\boldsymbol{\kappa}} \boldsymbol{G}_i\right) + \boldsymbol{\lambda} \boldsymbol{G}_i + \boldsymbol{\omega}(t - 2004)\boldsymbol{G}_i + \nu_t + \mu_{dgt} + \varepsilon_{it}, \tag{24}$$

where $\boldsymbol{\zeta}_t$ is a vector of time-varying coefficients and $\hat{\boldsymbol{\kappa}}\boldsymbol{G}_i$ is the fitted value from an estimate of Equation (20) using vaping in all years observed. Because we have demographic group intercepts and time trends and $\hat{\boldsymbol{\kappa}}\boldsymbol{G}_i$ varies only by demographic group, we must omit at least two years from the $\boldsymbol{\zeta}_t$ parameters. The more years we omit, the more precisely we can estimate the time trends $\boldsymbol{\omega}$. We estimate one indicator for the combined 2004–2010 period and one for each individual year after, omitting 2012, the year before vaping starts to become popular.

5.3 Event Study Figures

Figure 7 presents estimates of the $\boldsymbol{\zeta}_t$ parameters from Equation (24), the event study specification. For adults, the 2004–2010 and 2011 indicators are very close to the omitted year (2012), implying no differential smoking trends prior to e-cigarette introduction for demographic groups with higher versus lower e-cigarette demand. The estimates are not statistically distinguishable from zero in any year.

For youth, the 2004–2010 point estimate is below the omitted year, and the 2011 estimate is slightly above, although the latter difference is not statistically significant with 95 percent confidence. Consistent with Figure 6, the point estimates are very close to zero in the years after e-cigarettes are introduced.

NHIS and post-2011 BRFSS.

5.4 Parameter Estimates

Figure 8 presents separate estimates of Equation (21) for adults and youth. The first row of each panel presents our preferred estimates. For adults (youth), the primary point estimates are $\hat{\sigma}_\theta \approx 0.03$ ($\hat{\sigma}_\theta \approx 0.01$). This implies that groups that are ten percentage points more likely to vape on a given day reduced smoking by 0.003 (0.001) packs per day relative to trend. Both adult and youth estimates are statistically indistinguishable from zero. We can rule out σ coefficients of less than -0.16 or more than 0.29 for adults (less than -0.03 or more than 0.06 for youth) with 95 percent confidence.

The subsequent rows in each panel present robustness checks. *Control for 2003 smoking* allows the smoking trends to differ for demographics with higher versus lower initial smoking rates, by including an additional control for the 2003 smoking rate in person i's demographic cell and the interaction of that variable with a linear time trend. *Vaping begins in 2012* modifies the construction of $\tilde{q}^e_{it}$ in Equation (23) to use 2012 instead of 2013 as the year when e-cigarettes first saw non-negligible use. The standard errors widen slightly as the linear demographic time trends $\boldsymbol{\omega}$ must be estimated off fewer years, but the point estimates do not change much. *No imputed vaping data* uses only observed vaping q^e_{it} instead of imputing missing q^e_{it} beginning in 2013.

In the youth estimates, *Demog. cell predictors* uses demographic cells, rather than linear demographic groups, in $\boldsymbol{G}_i$. *Drop race other/missing* is motivated by Appendix Figure A4, which shows that the predicted vaping among people whose race is other/missing differs in MTF versus NYTS.

The next set of robustness checks, *Predictors excl. age (or race, etc.)* omit age (or race, or other demographic categories) from the vaping predictors $\boldsymbol{G}_i$. These are informal overidentification tests, allowing us to see whether the results are driven by any one demographic category. Consistent with the earlier informal overidentification tests in Appendix Figures A10–A13, the point estimates move little when we exclude any given demographic category. The standard errors illustrate that most of the identifying variation is from age (for adults) and grade (for youth), consistent with fact that these are the most predictive demographic categories illustrated in Figure 4.

The final set of robustness checks presents estimates using each dataset individually in the second stage regression. Our primary results from combining three datasets are about the average of the estimates from each individual dataset. The point estimates differ somewhat across datasets, which highlights the importance of our efforts to use all available data.

To argue that vaping is a material complement or substitute to smoking over our sample period, one would have to believe that some unobserved force increased or decreased smoking over the exact period that vaping became popular, breaking a previously steady downward trend. One would also have to believe that this unobserved force affected all demographic groups. And since the primary σ estimates are similar in the RMS data and the sample surveys, one would have to believe that this force could confound both sets of estimates.

Table 3 helps to put the primary results in context. We multiply $\hat{\sigma}_\theta$ for $\theta \in \{\text{adults}, \text{youth}\}$

by 2018 average vaping q^e_θ to estimate the change in smoking caused by the introduction of e-cigarettes. For the average adult, we can reject with 95 percent confidence that vaping increased (decreased) smoking by 0.007 (0.004) packs per day, or about 8 percent (4 percent) of average cigarette consumption. For the average youth, we can rule out with 95 percent confidence that vaping increased (decreased) smoking by more than 0.003 (0.001) packs per day, or about 51 (21) percent of average consumption. The percent terms are larger for youth because they already had low baseline smoking, but both substitution parameters are economically precise zeros in the sense that they rule out any material gateway effect (long-run complementarity) from vaping to smoking through 2018. We cannot rule out gateway effects that have not yet manifested themselves as of the 2018 surveys—for example, if high-vaping youth demographics will transition to smoking over a longer period.

Aggregating across all adults and youth, we can rule out that the introduction of e-cigarettes increased (decreased) smoking by more than about 660 (354) million packs in 2018. Furthermore, we can rule out that the introduction of e-cigarettes changed cigarette demand by more than 5 to 11 percent of the total decrease observed from 2004–2018. Thus, these estimates suggest that while e-cigarettes may be smoking cessation aids from some people and gateways to smoking for others, neither of these effects dominates in an economically significant way.

An important caveat to our welfare analysis is that we do not study the effects of e-cigarettes on other behaviors that may involve uninternalized harms. One key concern is that e-cigarettes may make it easier to consume marijuana. In Appendix D.2, we study trends in marijuana use among youth, for whom this concern is particularly salient. We find no evidence that youth marijuana use grew as e-cigarettes became popular. However, marijuana is increasingly being consumed through vaping, and this may be a more harmful form of consumption: the 2,807 lung injuries and 68 deaths from vaping in 2019 and early 2020 were primarily linked to marijuana e-liquids (Centers for Disease Control 2020).

6 Expert Survey

In this section, we present results from the E-cigarette Expert Survey. Figure 9 presents the distribution of responses to the first two questions on the survey: the ratios of internalities and externalities from daily vaping relative to daily smoking for the average adult. The mean (median) expert believes that the externalities and internalities from vaping are 48 (25) and 101 (80) percent as large as those from smoking cigarettes. There is substantial disagreement: the interquartile ranges of these two harm ratios are 10 to 75 percent and 27 to 180 percent, respectively.

The mean and median expert believe that the uninternalized harms from vaping are about twice as large for youth as they are for adults; see Appendix Figure A16 for the full distribution across experts. These three key results—material harms relative to combustible cigarettes, substantial uncertainty, and much larger harms for youth compared to adults—will be central to our welfare analysis.

Our sample includes both economists and health researchers cited in the National Academy of Sciences (2018) e-cigarette study. Appendix Figure A15 shows that these two subgroups report almost the exact same average internality ratio for vaping relative to smoking, while economists report slightly lower external harm and youth/adult harm ratios than health experts report. As we discuss below, the relative internality will be particularly important for our policy analysis, so it is notable that these two subgroups agree.

These results paint a substantially different picture than might have been expected on the basis of prior literature. A prominent early estimate suggested that e-cigarettes were only five percent as harmful as combustible cigarettes (Nutt et al. 2014). Viscusi (2016) argued that early evidence suggested that vaping could be at least 100 times safer than smoking. The National Academy of Sciences (2018, page 1) concluded that "e-cigarettes are likely to be far less harmful than combustible tobacco cigarettes." This difference is remarkable given that our experts are the same people who wrote the research cited in that study.

There are several potential reasons for these differences. First, while prior estimates quantified health harms, our expert survey focused on uninternalized externalities and internalities. This is crucial: our model from Section 1 clarifies that externalities and internalities are what matter for policy, not health harms *per se*, and these concepts are different. For example, some internalities might be caused by failing to anticipate the financial cost of addiction, and health harms can result in increased health care costs (a negative externality) or early mortality (a potential positive fiscal externality).

Second, our survey may reflect recent changes in the state of expert knowledge. Nutt et al. (2014) wrote that there was a "lack of hard evidence" for their conclusions, and Eissenberg et al. (2020) argue that e-cigarettes and e-liquids are more harmful than they were a few years ago and that "evidence of potential harm has accumulated."

Third, our results might suffer from common problems with unincentivized surveys, such as noisy or expressive responses. Fourth, our 33 percent response rate could suggest sample selection bias. However, even if all non-respondents would have rated vaping as zero percent as harmful as smoking, the average beliefs about relative harms in our full sampling frame would still be 33 percent as large as our estimates, and thus larger than Nutt et al.'s (2014) five percent.

Our final two questions elicited experts' views on public policy. About 52 percent of experts support a ban on flavored e-cigarettes, and 22 percent support a ban on all e-cigarettes. This disagreement among experts highlights the importance of quantitative policy analysis.

7 Optimal Regulation

7.1 Parameter Calibrations

In this section, we estimate the optimal e-cigarette tax using Equation (14) and the welfare effects of an e-cigarette ban using Equation (15). We use Monte Carlo simulations to capture the sampling variation in each parameter. Specifically, we re-estimate Equations (14) and (15) one million times, drawing each parameter in each equation from its distribution. Unless otherwise stated below, we draw each parameter from a normal distribution with mean and standard deviation equal to its point estimate and standard error.

Table 4 summarizes the parameters, their mean values in our primary simulations, and their sources. To further acknowledge uncertainty, we will also consider alternative assumptions for the key parameters in the next section. We use parameters from 2018, the most recent available year, and we inflate monetary amounts to 2018 dollars. We consider two consumer types $\theta \in \{a, y\}$, representing adults and youth.

We use the empirical estimate of η from Table 2 and the adult and youth σ from Figure 8. To avoid positive own-price elasticities, we re-draw any positive draw of η; this happens in only about 0.02 percent of simulations. We compute s_θ, the share of each type, by calculating the number of youth ages 12–17 and adults ages 18–100 in the 2018 American Community Survey.

Current youth and adult e-cigarette consumption $q^e_\theta(\tilde{\boldsymbol{p}})$ are the 2018 averages from the sample surveys plotted in Figure 2. Vaping is now in units of milliliters (ml) per person-day, and the e-cigarette tax rate and marginal distortion are in dollars per ml. We transform q^e_θ from the original survey units (share of days) to ml/person-day using Γ, the e-liquid consumption on an average vaping day from our E-cigarette User Survey.

We import the cigarette average marginal bias and externality from existing literature. For lack of data, we assume that these parameters are homogeneous across types.[21] We follow Cutler et al. (2015) in assuming that the marginal bias is $\gamma^c = (1-\beta)H_S$, where β is the present focus parameter and H_S is the discounted future private cost of smoking per pack. As we showed in an example in Section 1, this is the correct formula for marginal bias if present focus is the only behavioral bias, the social planner uses the long-run criterion (so that normative utility uses exponential discounting), and there is no habit formation. With habit formation, γ^c would be smaller with sophisticated present focus and probably larger with naive present focus (Gruber and Kőszegi 2001). Projection bias would probably increase γ^c. We use the stylized $\gamma^c = (1-\beta)H_S$ because of these modeling

[21] In reality, the cigarette distortion varies across people. For example, internalities and externalities could be larger for young people who are not yet addicted to nicotine. Furthermore, some correlation studies show that full smoking cessation is associated with disproportionately large health gains compared to partial smoking reduction (Song, Sung and Cho 2008) and that "dual use" of cigarettes and e-cigarettes is associated with particularly large toxicant exposures compared to only smoking or vaping (Goniewicz et al. 2018), although it is not clear whether these results are causal.

uncertainties.

We assume that the present discounted private health cost from smoking is $H_S = \$44.40$ per pack, inflating the estimate from Gruber and Kőszegi (2001) to 2018 dollars. We import $\beta = 0.67$ and its standard error from the Chaloupka, Levy and White (2019) experimental estimates of internalities from smoking. This gives a mean estimate of $\gamma^c = (1-\beta)H_S \approx \$14.65$. We assume that the marginal externality from smoking is $\phi^c \approx \$3.21$ per pack, inflating the estimate from Sloan et al. (2004) to 2018 dollars.[22] We assume that H_S and ϕ^c are known with certainty, as there are no standard errors in the original sources. Adding the internality and externality, our mean estimate of the smoking marginal distortion is $\varphi^c \approx \$17.86$ per pack.

To estimate the internalities and externalities from adult vaping, we multiply the existing estimates of smoking marginal bias and externality by the relative harms from vaping elicited from the E-cigarette Expert Survey. Define $\alpha_a^\gamma := \gamma_a^e/\gamma_a^c$ and $\alpha_a^\phi := \phi_a^e/\phi_a^c$ as the ratios of marginal internalities and externalities, respectively, for daily vaping versus daily smoking for adults. We use the means and empirical distributions of α_a^γ and α_a^ϕ from Figure 9, and we draw the tuple $(\alpha_a^\gamma, \alpha_a^\phi)$ jointly from the distribution of experts to allow these parameters to be correlated. We transform the α parameters from their original units (distortion per every-day vaper / distortion per every-day smoker) using Γ and $\Omega_a \approx 0.51$, the cigarette consumption on the average day smoked for adults in the sample surveys. The marginal distortion from adult vaping is then $\varphi_a^e = \frac{\Omega_a}{\Gamma}\left(\alpha_a^\gamma \gamma^c + \alpha_a^\phi \phi^c\right)$, in units of \$/ml.

To estimate the marginal distortion from youth vaping, we multiply the adult marginal distortion by the ratio of youth to adult harms elicited on the expert survey: $\varphi_y^e = \rho\varphi_a^e$. We draw ρ from the empirical distribution in Figure A16.

We use the 2018 population-weighted average tax rates $\tilde{\tau}^c$ and $\tilde{\tau}^e$ across states, and we use $\tilde{p}^e \approx \$3.90$ per ml from our E-cigarette User Survey. Note that only about 1/4 of the U.S. population lives in states, counties, or cities with e-cigarette taxes, so $\tilde{\tau}^e$ is about 1/4 of the average tax rate in areas that currently have taxes.

Appendix F provides additional details about empirical implementation.

7.2 Optimal Regulation Results

Three key parameters. Three key parameters drive our results on optimal regulation. First, we estimate that e-cigarette demand is more than unit elastic. This relatively elastic demand reduces the perceived consumer surplus from vaping, pushing toward the possibility that a ban might increase welfare.

Second, our point estimates of the substitution parameter σ imply very limited complementarity

[22] Most of this amount represents cross-subsidies from non-smokers to smokers that may no longer occur because most life insurance policies adjust for smoking status (DeCicca, Kenkel and Lovenheim 2020). On the other hand, health care costs have increased substantially in the past two decades.

or substitutability between e-cigarettes and cigarettes. This means that in our mean Monte Carlo simulation, optimal e-cigarette policy places little weight on cigarette market distortions. However, cigarette market distortions will matter for simulation draws with σ further from zero.

Third, prior research suggests large bias and externalities associated with smoking, and our experts believe that vaping is almost as harmful as smoking. This will imply that the optimal e-cigarette tax is much higher than current levels and that a ban might even increase welfare relative to current taxes. To see this, recall that the smoking marginal distortion is $\varphi^c \approx \$17.86$ per pack. Define $\alpha_a := \frac{\alpha_a^\gamma \gamma^c + \alpha_a^\phi \phi^c}{\gamma^c + \phi^c}$ as the relative uninternalized harms from vaping every day compared to smoking every day. At our mean parameter values, we have $\alpha_a \approx \frac{1.01 \cdot \$14.65 + 0.48 \cdot \$3.21}{\$14.65 + \$3.21} \approx 0.92$, implying that our average expert thinks that the marginal distortion from vaping every day is 92 percent as large as the marginal distortion from smoking every day. $\Omega_a \approx 0.51$ packs of cigarettes ($\Gamma \approx 0.58$ ml) are consumed in an average day of smoking (vaping). Thus, the adult vaping marginal distortion per milliliter is slightly less than the smoking marginal distortion per pack: $\varphi^e = \varphi^c \alpha_a \Omega_a / \Gamma \approx 17.86 \times 0.92 \times 0.51/0.58 \approx \$14.45/\text{ml}$. The optimal tax will be higher than that, because the youth vaping marginal distortion is higher by proportion $\rho \approx 2.11$. In comparison, the current average e-cigarette tax in states, counties, and cities that have taxes is \$0.89/ml, more than an order of magnitude less. One may disagree with our experts or with our other parameter assumptions, but it would take very different assumptions to change the conclusion that existing e-cigarette taxes are too low in the context of our model.

Optimal taxes. Figure 10 presents the distribution of optimal taxes over Monte Carlo simulation draws. As discussed in Section 1, the optimal tax could be negative (i.e. a subsidy) if cigarettes are much more harmful than e-cigarettes and the two goods are substitutes. Neither of these is the case at our mean estimates, and the optimal tax is positive in 97 percent of simulations. The vertical line marks the current average e-cigarette tax in states and local areas that have taxes, \$0.89/ml. The optimal tax exceeds that current average in about 95 percent of simulations. Thus, the model predicts with high confidence that it is optimal to impose some positive e-cigarette tax, and indeed a larger tax than the current norm.

Welfare effects of a ban. Figure 11 presents the welfare effects of an e-cigarette ban, separately for youth and adults. Recall that in our model, the optimal tax is always preferred to a ban, and we compare a ban to the status quo with current tax rates.

A successfully implemented youth vaping ban increases welfare in 92 percent of simulations, and the welfare gains in the median simulation are about \$181 per youth per year. There is a large mass of draws with welfare gains above \$200, driven by the right tail of the ρ distribution from experts who reported that vaping is very harmful for youth relative to adults.[23] While the U.S. already bans e-cigarette sales to minors, complete bans including sales to adults are rare and more controversial. The model predicts that an adult vaping ban would increase welfare in about 90

[23]Appendix Figure A17 presents parameter regions where the youth ban increases welfare.

percent of simulations, and the welfare gains in the median simulation are about \$50 per American adult per year.

Sources of uncertainty. What parameters generate the most uncertainty in setting optimal policy? Figure 12 presents the variance in predicted welfare effects of a ban from Monte Carlo simulations that hold each listed parameter fixed, as a fraction of the variance in the primary simulations in Figure 11. The key message is that the α_a parameters, and in particular α_a^γ, contribute by far the most to policy uncertainty. Tangibly, our e-cigarette experts disagree substantially about the harms from vaping compared to smoking, and this disagreement matters for policy evaluation. An equivalent figure showing the contribution of each parameter to uncertainty in the optimal tax tells the same story.

Welfare effects at different α_a. Given that α_a is the key source of policy uncertainty, Figure 13 presents the mean and 95 percent confidence intervals for the welfare effects of an e-cigarette ban for a range of α_a from 0 to 2. Note that because η enters Equation (15) in the denominator, the bounds on ΔW are not symmetric; for draws of η close to 0, the perceived consumer surplus from e-cigarettes becomes large.

Nutt et al. (2014) concluded that e-cigarettes were about 5 percent as harmful as combustible cigarettes. At $\alpha_a = 0.05$, a ban has almost exactly zero welfare effect in the mean simulation. At about $\alpha_a \approx 1/3$, meaning that vaping has about 1/3 the uninternalized harms of smoking, the model predicts with 95 percent confidence that a ban increases welfare relative relative to current tax rates.

Alternative assumptions. To further understand the sources of policy uncertainty, Table 5 presents optimal tax rates and welfare effects of a ban under alternative assumptions. In each row of Panels (a) and (b), we present the mean τ^{e*} or $\Delta\overline{W}$ at the parameter assumption listed in the first column for the 25th percentile, mean, and 75th percentile of α_a across our experts, drawing the other parameters from their distributions. The first 12 rows are parallel across the two panels.

Row 1 presents the primary specification. At the mean parameter values, the optimal tax is \$18/ml. Because this is so much higher than the current average e-cigarette tax $\tilde{\tau}^e$, banning e-cigarettes is preferred to the status quo. A complete e-cigarette ban for youth and adults increases welfare by \$91 per person per year, or \$25 billion per year over the 279 million people aged 12 and older nationwide. Even at the 25th percentile of experts' α_a distribution, the optimal e-cigarette tax is considerably larger than the average existing state or local e-cigarette tax, and a ban increases welfare.

Rows 2–5 present alternative assumptions for the substitution parameter σ. Since smoking generates uninternalized distortions ($\varphi^c > \tilde{\tau}^c$), more substitutability (more negative σ) pushes toward a lower optimal tax and lower welfare gains from a ban, and more complementarity (more positive σ) pushes in the other direction. Row 2 uses the minimum distance estimates from the Nielsen RMS data in Section 4, $\hat{\sigma}_{\text{adult}} \approx -0.056$ and $\sigma_{\text{youth}} \approx -0.000$, which suggest slightly more substitutability between smoking and vaping. This pushes toward a lower optimal tax and lower welfare gains from a ban. Row 3 uses the $\hat{\sigma}$ parameters from both Sections 4 and 5, combined using a minimum distance estimator as described in Appendix D.1.

Rows 4 and 5 present results assuming that an average day of vaping is a perfect complement ($\sigma = 0.5$ for adults and $\sigma = 0.15$ for youth) or a perfect substitute ($\sigma = -0.5$ for adults and $\sigma = -0.15$ for youth) for an average day of smoking. For perfect substitutes and lower values of α_a, it is optimal to subsidize e-cigarettes, and a ban reduces welfare. These parameters are consistent with some policy arguments to encourage e-cigarettes as a harm-reduction approach for existing smokers, but only if α_a is on the low end of our experts' beliefs.

Row 6 allows the social planner to set the optimal cigarette tax $\tau^c = \tau^{c*} = \varphi^c$. Optimal policy then considers the e-cigarette market in isolation. Because our primary estimates already use $\sigma \approx 0$, this makes little difference relative to Row 1.

Rows 7–9 present alternative assumptions for the cigarette marginal bias γ^c and externality ϕ^c. This significantly affects optimal policy, primarily because the e-cigarette marginal distortions φ^e_θ also depend on these parameters. Row 7 assumes that the conventional wisdom of policymakers is more informative about the marginal distortion from smoking than the academic research we use in our primary estimates. In that row, we assume that existing average cigarette taxes $\tilde{\tau}^c$ are set optimally and rescale γ^c and ϕ^c so that $\varphi^c = \tilde{\tau}^c$. This substantially reduces the resulting e-cigarette marginal distortions φ^e_θ, reducing the optimal e-cigarette tax and suggesting that a ban reduces welfare at lower α_a values.

Rows 8 and 9 set the cigarette internality γ^c at \$1.80 per pack and \$80 per pack, to match the estimates in Gruber and Kőszegi (2001) and Chaloupka, Levy and White (2019), respectively. Even in row 8, banning e-cigarettes increases welfare, and the optimal tax is considerably higher than the current norm at the mean α_a.

Rows 10 and 11 consider the youth and adult markets in isolation. Since the mean $\rho \approx 2.11$, the optimal e-cigarette tax for youth is about twice as large as the optimal tax for adults. The per-youth gains from a youth-specific ban are 4–5 times larger than the per-adult gains from an adult-specific ban. These calculations underscore that if leakage or enforcement issues make it easier to impose type-specific bans than type-specific taxes, there are parameter configurations under which a youth sales ban plus a tax on the remaining sales to adults could be the constrained optimum in our model.

Row 12 further highlights the uncertainty in our model's policy prescriptions. For this row, we use the $\hat{\sigma}$ and $\hat{\eta}$ from column 5 of Table 2, the estimates without cluster-specific linear time trends, which imply more substitutability and more inelastic demand. The more negative $\hat{\sigma}$ implies a substitution distortion of about \$-5/ml, meaning that it is optimal to subsidize e-cigarettes to reduce uninternalized harms from cigarette use. We also use the argument of Viscusi (2016) that $\gamma^e < 0$ (i.e. bias causes people to vape too little) because people underestimate health risks. Viscusi

(2016, Table 2) finds that people believe that the mortality risks from typical vaping are 66 percent as large as those from smoking a pack of cigarettes per day (total mortality of 33.3 versus 50.3 out of 100). He cites other research to argue that the true relative risk is only 1 percent. Thus, e-cigarette users overestimate the health harms of e-cigarettes by $(0.66 - 0.01)H_S \approx 0.65 \cdot \$44.40 \approx \$29$ per day of vaping. Translating to \$/ml using Γ and assuming that this is the only e-cigarette distortion implies $\varphi^e \approx$ \$-51/ml. The results in Row 12 show that under these assumptions, the optimal policy is to heavily subsidize e-cigarettes instead of taxing or banning them.

Panel (b) of Table 5 includes additional rows that are particularly relevant for evaluating a ban. Rows 13 and 14 present alternative assumptions for the elasticity η. Since we assume that the perceived consumer surplus loss from a ban is the area under a line drawn tangent to the demand curve at current prices, more inelastic demand implies larger perceived consumer surplus loss. The resulting changes in perceived consumer surplus are still far outweighed by the uninternalized distortions.

In many markets, youth are more price elastic than adults, but our aggregate Nielsen RMS data do not allow us to estimate separate elasticities by type. Row 15 keeps the adult elasticity at the RMS estimate and assumes that youth demand is 50 percent more elastic. This increases the net welfare gains from a ban, but the change relative to Row 1 is small because youth represent only nine percent of the population.

The e-liquid price in Nielsen RMS is $\tilde{p}^e \approx \$6.58$/ml, somewhat higher than in our E-cigarette User Survey. Using this higher value in Row 16 increases the perceived consumer surplus loss, which reduces the net welfare gains from a ban.

Rows 17 and 18 present alternative assumptions for the baseline e-cigarette tax $\tilde{\tau}^e$. Row 17 assumes that $\tilde{\tau}^e = \varphi^e$, so there is no uninternalized e-cigarette distortion. Since $\sigma \approx 0$, setting the optimal e-cigarette tax removes essentially all economic justification for a ban, and the welfare gain from a ban equals the perceived consumer surplus change, which is negative. However, current tax rates are far below our primary estimate of φ^e. Row 18 shows that even doubling $\tilde{\tau}^e$ still leaves ample welfare gains from a ban. Thus, any factors that make it difficult to materially raise e-cigarette taxes could provide justification for a ban if our experts are right about the uninternalized harms.

Row 19 presents a set of combined assumptions that reduce the welfare gains from a ban: reducing the demand elasticity to $\eta = -0.5$ (increasing perceived consumer surplus), halving the cigarette marginal distortion (and thus halving the resulting e-cigarette distortion), and making vaping a slight substitute (instead of slight complement) for smoking. Under these assumptions, a ban still increases welfare at the mean α_a, but it reduces welfare at lower draws of α_a.

8 Conclusion

Electronic cigarettes are one of the most controversial new products of the past decade, due to uncertainty about their health effects and whether they are primarily a cessation aid or a gateway drug for traditional combustible cigarettes. We lay out a simple dynamic behavioral optimal policy framework that delivers formulas for the optimal e-cigarette tax and welfare effects of a ban as functions of several sufficient statistics. We estimate these statistics using Nielsen RMS scanner data, sample surveys, and a new survey of e-cigarette experts. We find that e-cigarette demand is price elastic, vaping is neither a significant complement nor substitute for smoking combustible cigarettes over the medium term, and our experts believe that vaping is almost as harmful as smoking.

Our Monte Carlo simulations make clear that parameter uncertainty generates substantial uncertainty in policy predictions, even within the context of our model. However, three conclusions seem robust. First, since most of the policy uncertainty in our model is driven by uncertainty over the uninternalized externalities and internalities from vaping, more research on those parameters would be very valuable. Second, eliminating youth vaping increases welfare in 92 percent of our model simulations, suggesting that existing regulations banning e-cigarette sales to minors and all sales of flavored e-cigarettes (which are especially appealing to minors) probably increase welfare. Third, if our experts are correct about the large uninternalized harms from vaping, in our model the optimal tax on e-cigarettes is probably much higher than the current norm.

References

Abouk, Rahi and Scott Adams, "Bans on Electronic Cigarette Sales to Minors and Smoking Among High School Students," *Journal of Health Economics*, 2017, *54*, 17–24.

_ , _ , Bo Feng, Johanna Catherine Maclean, and Michael F Pesko, "The Effect of E-Cigarette Taxes on Pre-Pregnancy and Prenatal Smoking, and Birth Outcomes," 2019. NBER Working Paper No. 26126.

Aguiar, Luis and Joel Waldfogel, "Quality Predictability and the Welfare Benefits From New Products: Evidence From the Digitization of Recorded Music," *Journal of Political Economy*, 2018, *126* (2), 492–524.

Allcott, Hunt and Dmitry Taubinsky, "Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, 2015, *105* (8), 2501–2538.

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

Ashley, Elizabeth M, Clark Nardinelli, and Rosemarie A Lavaty, "Estimating the Benefits of Public Health Policies That Reduce Harmful Consumption," *Health Economics*, 2015, *24* (5), 617–624.

Bartik, Timothy J, *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research, 1991.

Bernheim, B. Douglas and Antonio Rangel, "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," *Quarterly Journal of Economics*, 2009, *124* (1), 51–104.

_ and Dmitry Taubinsky, "Behavioral Public Economics," in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *The Handbook of Behavioral Economics*, Vol. 1, New York: Elsevier, 2018.

Berry, Kaitlyn M, Jessica L Fetterman, Emelia J Benjamin, Aruni Bhatnagar, Jessica L Barrington-Trimis, Adam M Leventhal, and Andrew Stokes, "Association of Electronic Cigarette Use With Subsequent Initiation of Tobacco Cigarettes in US Youths," *JAMA Network Open*, 2019, *2* (2), e187794–e187794.

Blanchard, Olivier Jean and Lawrence F Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *1*, 1–75.

Bour, Norm, "Vaping Trends: Early Adopters and Future Directions," https://vapementors.com/vaping-trend-early-adapters-and-future-directions/ 2019. Vape Mentors.

Boxell, Levi, Matthew Gentzkow, and Jesse M Shapiro, "Greater Internet Use is Not Associated With Faster Growth in Political Polarization Among US Demographic Groups," *Proceedings of the National Academy of Sciences*, 2017, *114* (40), 10612–10617.

Cantrell, Jennifer, Jidong Huang, Marisa S Greenberg, Haijuan Xiao, Elizabeth C Hair, and Donna Vallone, "Impact of E-Cigarette and Cigarette Prices on Youth and Young Adult E-Cigarette and Cigarette Behaviour: Evidence from a National Longitudinal Cohort," *Tobacco Control*, 2019. https://tobaccocontrol.bmj.com/content/early/2019/06/05/tobaccocontrol-2018-054764.

Centers for Disease Control and Prevention, "Outbreak of Lung Injury Associated with E-Cigarette Use, or Vaping," https://www.cdc.gov/tobacco/basic_information/e-cigarettes/severe-lung-disease.html 2020.

Chaloupka, Frank J, Kenneth E Warner, Daron Acemoğlu, Jonathan Gruber, Fritz Laux, Wendy Max, Joseph Newhouse, Thomas Schelling, and Jody Sindelar, "An Evaluation of the FDA's Analysis of the Costs and Benefits of the Graphic Warning Label," *Tobacco Control*, 2015, *24* (2), 112–119.

_ **, Matthew R Levy, and Justin S White**, "Estimating Biases in Smoking Cessation: Evidence from a Field Experiment," 2019. NBER Working Paper No. 26522.

Chatterjee, Kshitij, Bashar Alzghoul, Ayoub Innabi, and Nikhil Meena, "Is Vaping a Gateway to Smoking: a Review of the Longitudinal Studies," *International Journal of Adolescent Medicine and Health*, 2016, *30* (3).

Chetty, Raj, "Behavioral Economics and Public Policy: A Pragmatic Perspective," *American Economic Review*, 2015, *105* (5), 1–33.

_ **, Adam Looney, and Kory Kroft**, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

Cooper, Michael T and Michael F Pesko, "The Effect of E-Cigarette Indoor Vaping Restrictions on Adult Prenatal Smoking and Birth Outcomes," *Journal of Health Economics*, 2017, *56*, 178–190.

Cotti, Chad, Charles Courtemanche, Catherine Maclean, Erik Nesson, Michael Pesko, and Nathan Tefft, "The Effects of E-Cigarette Taxes on E-Cigarette Prices and Tobacco Product Sales: Evidence from Retail Panel Data," 2020. NBER Working Paper No. 26724.

__ , Erik Nesson, and Nathan Tefft, "The Relationship Between Cigarettes and Electronic Cigarettes: Evidence from Household Panel Data," *Journal of Health Economics*, 2018, *61*, 205–219.

Cutler, David, Amber Jessup, Donald Kenkel, and Martha Starr, "Valuing Regulations Affecting Addictive or Habitual Goods," *Journal of Benefit Cost Analysis*, 2015, *6* (2), 247–280.

__ , __ , __ , and __ , "Economic Aproaches to Estimating Benefits of Regulations Affecting Addictive Goods," *American Journal of Preventive Medicine*, 2016, *50* (5), 520–526.

Dave, Dhaval, Bo Feng, and Michael F Pesko, "The Effects of E-Cigarette Minimum Legal Sale Age Laws on Youth Substance Use," *Health Economics*, 2019, *28* (3), 419–436.

DeCicca, Philip, Donald Kenkel, Feng Liu, and Hua Wang, "Behavioral Welfare Economics and FDA Tobacco Regulations," in "Human Capital and Health Behavior," Emerald Publishing Limited, 2017, pp. 143–179.

__ , Donald S Kenkel, and Michael F Lovenheim, "The Economics of Tobacco Regulation: A Comprehensive Review," 2020. NBER Working Paper No. 26923.

Diamond, Peter A., "Consumption Externalities and Imperfect Corrective Pricing," *The Bell Journal of Economics and Management Science*, 1973, *4* (2), 526–538.

Eissenberg, Thomas, Aruni Bhatnagar, Simon Chapman, Sven-Eric Jordt, Alan Shihadeh, and Eric K. Soule, "Invalidity of an Oft-Cited Estimate of the Relative Harms of Electronic Cigarettes," *American Journal of Public Health*, 2020, *110* (2), 161–162.

Farhi, Emmanuel and Xavier Gabaix, "Optimal Taxation with Behavioral Agents," *American Economic Review*, January 2020, *110* (1), 298–336.

Friedman, Abigail S, "How Does Electronic Cigarette Access Affect Adolescent Smoking?," *Journal of Health Economics*, 2015, *44*, 300–308.

Gentzkow, Matthew, "Valuing New Goods in a Model With Complementarity: Online Newspapers," *American Economic Review*, 2007, *97* (3), 713–744.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," 2019. NBER Working Paper No. 24408.

Goniewicz, Maciej L, Danielle M Smith, Kathryn C Edwards, Benjamin C Blount, Kathleen L Caldwell, Jun Feng, Lanqing Wang, Carol Christensen, Bridget Ambrose, Nicolette Borek, Dana van Bemmel, Karen Konkel, Gladys Erives, Cassandra A Stanton, Elizabeth Lambert, Heather L Kimmel, Dorothy Hatsukami,

Stephen S Hecht, Raymond S Niaura, Mark Travers, Charles Lawrence, and Andrew J Hyland, "Comparison of Nicotine and Toxicant Exposure in Users of Electronic Cigarettes and Combustible Cigarettes," *JAMA Network Open*, December 2018, *1* (8), e185937.

Goolsbee, Austan and Amil Petrin, "The Consumer Gains from Direct Broadcast Satellites and the Competition With Cable TV," *Econometrica*, 2004, *72* (2), 351–381.

Gottlieb, Scott, "Statement from FDA Commissioner Scott Gottlieb, M.D., on New Steps to Address Epidemic of Youth E-Cigarette Use," https://www.fda.gov/news-events/press-announcements/statement-fda-commissioner-scott-gottlieb-md-new-steps-address-epidemic-youth-e-cigarette-use Sep 2018. U.S. Food and Drug Administration.

Gotts, Jeffrey E, Sven-Eric Jordt, Rob McConnell, and Robert Tarran, "What Are the Respiratory Effects of E-Cigarettes?," *BMJ*, 2019, *366.*

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

_ **and** _ , "Tax Incidence When Individuals Are Time-Inconsistent: The Case of Cigarette Excise Taxes," *Journal of Public Economics*, 2004, *88* (9-10), 1959–1987.

_ **and Sendhil Mullainathan**, "Do Cigarette Taxes Make Smokers Happier?," *The B.E. Journal of Economic Analysis & Policy*, 2005, *5* (1).

Hainmueller, Jens, "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 2012, *20* (1), 25–46.

Hajek, Peter, Anna Phillips-Waller, Dunja Przulj, Francesca Pesola, Katie Myers Smith, Natalie Bisal, Jinshuo Li, Steve Parrott, Peter Sasieni, Lynne Dawkins et al., "A Randomized Trial of E-Cigarettes Versus Nicotine-Replacement Therapy," *New England Journal of Medicine*, 2019, *380*, 629–637.

Hartwell, Greg, Sian Thomas, Matt Egan, Anna Gilmore, and Mark Petticrew, "E-Cigarettes and Equity: a Systematic Review of Differences in Awareness and Use Between Sociodemographic Groups," *Tobacco Control*, 2017, *26* (e2), e85–e91.

Hausman, Jerry, "Valuation of New Goods Under Perfect and Imperfect Competition," in Timothy Bresnaham and Roger Gordon, eds., *The Economics of New Goods*, Vol. 58, University of Chicago Press, 1996, pp. 207–248.

Huang, Jidong, Cezary Gwarnicki, Xin Xu, Ralph S Caraballo, Roy Wada, and Frank J Chaloupka, "A Comprehensive Examination of Own-and Cross-Price Elasticities of Tobacco and Nicotine Replacement Products in the US," *Preventive Medicine*, 2018, *117*, 107–114.

Jin, Lawrence, Don Kenkel, Feng Liu, and Hua Wang, "Retrospective and Prospective Benefit-Cost Analyses of US Anti-Smoking Policies," *Journal of Benefit-Cost Analysis*, 2015, *6* (1), 154–186.

Kenkel, Donald S, Sida Peng, Michael F Pesko, and Hua Wang, "Mostly Harmless Regulation? Electronic Cigarettes, Public Policy and Consumer Welfare," 2019. NBER Working Paper No. 23710.

Leventhal, Adam M, David R Strong, Matthew G Kirkpatrick, Jennifer B Unger, Steve Sussman, Nathaniel R Riggs, Matthew D Stone, Rubin Khoddam, Jonathan M Samet, and Janet Audrain-McGovern, "Association of Electronic Cigarette Use with Initiation of Combustible Tobacco Product Smoking in Early Adolescence," *JAMA*, 2015, *314* (7), 700–707.

Levy, David T, Kenneth E Warner, K Michael Cummings, David Hammond, Charlene Kuo, Geoffrey T Fong, James F Thrasher, Maciej Lukasz Goniewicz, and Ron Borland, "Examining the Relationship of Vaping to Smoking Initiation among US Youth and Young Adults: a Reality Check," *Tobacco Control*, 2019, *28* (6), 629–635.

Levy, Helen G, Edward C Norton, and Jeffrey A Smith, "Tobacco Regulation and Cost-Benefit Analysis: How Should We Value Foregone Consumer Surplus?," *American Journal of Health Economics*, 2018, *4* (1), 1–25.

Liber, Alex C and Kenneth E Warner, "Has Underreporting of Cigarette Consumption Changed Over Time? Estimates Derived From US National Health Surveillance Systems Between 1965 and 2015," *American Journal of Epidemiology*, 2018, *187* (1), 113–119.

Mullainathan, Sendhil, Joshua Schwartzstein, and William J Congdon, "A Reduced-Form Approach to Behavioral Public Finance," *Annual Review of Economics*, 2012, *4* (1), 511–540.

National Academies of Sciences, Engineering, and Medicine, *Public Health Consequences of E-Cigarettes*, National Academies Press, 2018.

Nevo, Aviv, "New Products, Quality Changes, and Welfare Measures Computed From Estimated Demand Systems," *Review of Economics and Statistics*, 2003, *85* (2), 266–275.

Nutt, David J, Lawrence D Phillips, David Balfour, H Valerie Curran, Martin Dockrell, Jonathan Foulds, Karl Fagerstrom, Kgosi Letlape, Anders Milton, Riccardo Polosa et al., "Estimating the Harms of Nicotine-Containing Products Using the MCDA Approch," *European Addiction Research*, 2014, *20* (5), 218–225.

O'Donoghue, Ted and Matthew Rabin, "Optimal Sin Taxes," *Journal of Public Economics*, 2006, *90* (10-11), 1825–1849.

Pepper, Jessica K, Sherry L Emery, Kurt M Ribisl, and Noel T Brewer, "How US Adults Find Out About Electronic Cigarettes: Implications for Public Health Messages," *Nicotine & Tobacco Research*, 2014, *16* (8), 1140–1144.

Perikleous, Evanthia P, Paschalis Steiropoulos, Emmanouil Paraskakis, Theodoros C Constantinidis, and Evangelia Nena, "E-Cigarette Use Among Adolescents: an Overview of the Literature and Future Perspectives," *Frontiers in Public Health*, 2018, *6*, 86.

Pesko, Michael F and Casey Warman, "The Effect of Prices on Youth Cigarette and E-Cigarette Use: Economic Substitutes or Complements?," *SSRN*, 2017. https://dx.doi.org/10.2139/ssrn.3077468.

_ **and Janet M Currie**, "E-Cigarette Minimum Legal Sale Age Laws and Traditional Cigarette Use Among Rural Pregnant Teenagers," *Journal of Health Economics*, 2019, *66*, 71–90.

_ **, Charles J Courtemanche, and Johanna Catherine Maclean**, "The Effects of Traditional Cigarette and E-Cigarette Taxes on Adult Tobacco Product Use," 2019. NBER Working Paper No. 26017.

_ **, Jenna M Hughes, and Fatima S Faisal**, "The Influence of Electronic Cigarette Age Purchasing Restrictions on Adolescent Tobacco and Marijuana Use," *Preventive Medicine*, 2016, *87*, 207–212.

_ **, Jidong Huang, Lloyd D Johnston, and Frank J Chaloupka**, "E-Cigarette Price Sensitivity Among Middle-and High-School Students: Evidence from Monitoring the Future," *Addiction*, 2018, *113* (5), 896–906.

Petrin, Amil, "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 2002, *110* (4), 705–729.

Pierannunzi, Carol, Machell Town, William Garvin, Frederick E. Shaw, and Lina Balluz, "Methodologic Changes in the Behavioral Risk Factor Surveillance System in 2011 and Potential Effects on Prevalence Estimates.," *MMWR. Morbidity and Mortality Weekly Report*, 2012, *61* (22), 410.

Primack, Brian A, Samir Soneji, Michael Stoolmiller, Michael J Fine, and James D Sargent, "Progression to Traditional Cigarette Smoking After Electronic Cigarette Use Among US Adolescents and Young Adults," *JAMA Pediatrics*, 2015, *169* (11), 1018–1023.

Rees-Jones, Alex and Kyle Rozema, "Price Isn't Everything: Behavioral Response around Changes in Sin Taxes," 2020. NBER Working Paper No. 25958.

Saffer, Henry, Daniel Dench, Dhaval Dave, and Michael Grossman, "E-Cigarettes and Adult Smoking," 2018. NBER Working Paper No. 24212.

_ , _ , **Michael Grossman, and Dhaval Dave**, "E-Cigarettes and Adult Smoking: Evidence from Minnesota," 2019. NBER Working Paper No. 26589.

Sloan, Frank A, Jan Ostermann, Donald H Taylor Jr, Christopher Conover, and Gabriel Picone, *The Price of Smoking*, MIT Press, 2004.

Soneji, Samir, Jessica L Barrington-Trimis, Thomas A Wills, Adam M Leventhal, Jennifer B Unger, Laura A Gibson, JaeWon Yang, Brian A Primack, Judy A Andrews, Richard A Miech et al., "Association Between Initial Use of E-Cigarettes and Subsequent Cigarette Smoking Among Adolescents and Young Adults: a Systematic Review and Meta-Analysis," *JAMA Pediatrics*, 2017, *171* (8), 788–797.

Song, Yun-Mi, Joohon Sung, and Hong-Jun Cho, "Reduction and Cessation of Cigarette Smoking and Risk of Cancer: A Cohort Study of Korean Men," *Journal of Clinical Oncology*, 2008, *26* (31), 5101–5106.

Stoklosa, Michal, Jeffrey Drope, and Frank J Chaloupka, "Prices and E-Cigarette Demand: Evidence from the European Union," *Nicotine & Tobacco Research*, 2016, *18* (10), 1973–1980.

Tax Foundation, "Vapor Taxes by State, 2018," https://taxfoundation.org/vapor-taxes-2018// 2018.

_ , "How High Are Vapor Excise Taxes in Your State?," https://taxfoundation.org/state-vapor-taxes-2019/ 2019.

Trajtenberg, Manuel, "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners," *Journal of Political Economy*, 1989, *97* (2), 444–479.

Viscusi, W Kip, "Do Smokers Underestimate Risks?," *Journal of Political Economy*, 1990, *98*, 1253–1269.

_ , "Risk Beliefs and Preferences for E-Cigarettes," *American Journal of Health Economics*, 2016, *2* (2), 213–240.

Watkins, Shannon Lea, Stanton A Glantz, and Benjamin W Chaffee, "Association of Noncigarette Tobacco Product Use With Future Cigarette Smoking Among Youth in the Population Assessment of Tobacco and Health (PATH) Study, 2013-2015," *JAMA Pediatrics*, 2018, *172* (11), 181–187.

Zeller, Mitch, "Response to the Epidemic of E-Cigarette Use, Especially Among Children," https://www.fda.gov/news-events/congressional-testimony/federal-response-epidemic-e-cigarette-use-especially-among-children-and-food-and-drug Dec 2019. U.S. Food and Drug Administration.

Zheng, Yuqing, Chen Zhen, Daniel Dench, and James M Nonnemaker, "US Demand for Tobacco Products in a System Framework," *Health Economics*, 2017, *26*, 1067–1086.

Table 1: **Smoking and Vaping Sample Surveys**

Dataset	**Population**	**Observations**	**Years**	**Notes**
BRFSS	Adults	5,346,115	2004–2018	Sampling change in 2011
MTF	Youth	591,740	2005–2018	Inconsistent race data in 2004
NHIS	Adults	412,888	2004–2018	
NSDUH	Adult sample	590,303	2004–2018	No vaping data
NSDUH	Youth sample	268,676	2004–2018	No vaping data
NYTS	Youth	227,813	2004, 2006, 2009, 2011–2018	

Notes: Datasets are the Behavioral Risk Factor Surveillance System (BRFSS), the National Health Interview Survey (NHIS), the National Survey of Drug Use and Health (NSDUH), Monitoring the Future (MTF), and the National Youth Tobacco Survey (NYTS)

Table 2: **Own- and Cross-Price Elasticity of Demand for E-cigarettes**

(a) **First Stage and Reduced Form**

	(1)	(2)	(3)
Dependent variable:	ln(e-cig price)	ln(cig price)	ln(e-cig units)
ln(e-cig % tax rate + 1)	0.580	0.126	-0.730
	(0.048)	(0.051)	(0.144)
ln(cig % tax rate + 1)	-0.011	0.620	0.181
	(0.061)	(0.087)	(0.285)
Observations	285,985	285,985	285,985

(b) **Instrumental Variables Estimates**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	ln(e-cig units)	ln(e-cig units)	ln(e-cig units)	ln(e-cig units)	ln(e-cig units)	ln(e-cig units)	ln(e-cig units)
ln(e-cig price)	-1.318	-1.513	-1.060	-0.924	-1.000	-1.406	-1.290
	(0.366)	(0.322)	(0.420)	(0.370)	(0.246)	(0.308)	(0.498)
ln(cig price)	0.267	0.612	0.588	0.610	0.655	0.327	0.447
	(0.459)	(0.550)	(0.538)	(0.527)	(0.344)	(0.458)	(0.599)
UPC-cluster FE	Yes	No	Yes	Yes	Yes	Yes	Yes
UPC-month FE	Yes	No	No	Yes	Yes	Yes	Yes
Division-month FE	Yes	No	No	No	Yes	Yes	Yes
Cluster × month trend	Yes	No	No	No	No	Yes	Yes
Quasi-panel	No	No	No	No	No	No	Yes
Time-varying state controls	Yes	Yes	Yes	Yes	Yes	No	Yes
Observations	285,985	286,491	286,303	285,985	285,985	285,985	499,664

Notes: This table presents estimates of the own- and cross-price elasticity of demand for e-cigarettes from Equation (16), using UPC-cluster-month data. There are 51 geographic clusters: the two counties that have e-cigarette taxes, each of the contiguous 48 states (excluding those two counties), and Washington, D.C. Standard errors are clustered by geographic cluster. Observations are weighted by the UPC's sales in non-taxed clusters in that calendar year, divided by total sales across all UPCs in that year in non-taxed clusters. Panel (a) presents the first stage and reduced form, using the same set of controls as in our primary estimate in column 1 of Panel (b). Panel (b) presents the instrumental variables estimates. Time-varying state controls are the state unemployment rate and beer tax rate as well as indicators for whether the state has an indoor vaping ban, has a medical marijuana law, has passed or implemented a prescription drug program, and implemented the Medicaid expansion. Column 7 presents estimates in a "quasi-panel" in which we add zero-sales observations for all UPCs that had non-zero sales in cluster s in any prior month, beginning with the month in which the UPC first had sales.

Table 3: **Effects of Vaping on Smoking**

	Adults	**Youth**
$\hat{\sigma}$ (packs per day/share of days)	0.03	0.01
95% confidence interval	(-0.16, 0.29)	(-0.03, 0.06)
2018 average vaping (share of days)	0.024	0.053
Effect of vaping on smoking (packs/day)	0.00083	0.00068
95% confidence interval	(-0.00374, 0.00690)	(-0.00138, 0.00329)
2018 average smoking (packs/day)	0.082	0.006
Effect of vaping on smoking (%)	1.0	10.6
95% confidence interval	(-4.5, 8.4)	(-21.4, 51.2)
2018 implied total smoking (million packs)	7,495	58.7
Effect of vaping on smoking (million packs)	76.0	6.2
95% confidence interval	(-340.9, 629.7)	(-12.6, 30.0)
2004–2018 smoking decrease (packs/day)	0.071	0.030
Effect of vaping on smoking (% of decrease)	-1.2	-2.3
95% confidence interval	(-9.8, 5.3)	(-11.1, 4.7)

Notes: This table presents estimates of the substitution parameter $\sigma_\theta := \frac{dq_\theta^c}{dq_\theta^e}$ and further analysis. We compute the effect of vaping on smoking (packs/day) by multiplying $\hat{\sigma}$ by average vaping. We compute the effect of vaping on smoking (%) by dividing the effect of vaping on smoking (packs/day) by average packs per day smoked in 2018. We compute the effect of vaping on smoking in 2018 (million packs) by multiplying the effect of vaping on smoking (%) by the total smoking in 2018 (million packs) implied by the sample survey data. We compute the effect of vaping on smoking (% of decrease) by dividing the effect of vaping on smoking (packs per day) by the change in packs per day smoked from 2004–2018. The confidence intervals for $\hat{\sigma}$ reflect the 2.5th and 97.5th percentiles of estimates from 200 bootstrap replications, where we draw bootstrap samples by demographic cell.

Table 4: **Parameters for Policy Analysis**

Object	**Description and units**	**Mean**	**Data source**
η	E-cigarette own-price elasticity	-1.318	RMS (Table 2)
σ_{adult}	E-cig effect on smoking (packs/day vaped)	0.035	Figure 8
σ_{youth}	E-cig effect on smoking (packs/day vaped)	0.013	Figure 8
s_{adult}	Population share adults	0.910	2018 American Community Survey
s_{youth}	Population share youth	0.090	2018 American Community Survey
q^e_{adult}	Share of person-days vaped	0.024	BRFSS, NHIS 2018
q^e_{youth}	Share of person-days vaped	0.053	MTF, NYTS 2018
Γ	Average e-liquid use (ml/day vaped)	0.58	E-cigarette User Survey
H_S	Present discounted cost of smoking per pack	44.4	Gruber and Kőszegi (2001)
β	Present bias	0.670	Chaloupka et al. (2019)
ϕ^c	Cigarette marginal externality ($/pack)	3.21	Sloan et al. (2004)
$\alpha^{\gamma}_{\text{adult}}$	$\frac{\text{Internality from vaping every day}}{\text{Internality from smoking every day}}$	1.015	E-cigarette Expert Survey
$\alpha^{\phi}_{\text{adult}}$	$\frac{\text{Externality from vaping every day}}{\text{Externality from smoking every day}}$	0.481	E-cigarette Expert Survey
ρ	Ratio of youth to adult e-cigarette distortion	2.112	E-cigarette Expert Survey
Ω_{adult}	Average cigarette use (packs/day smoked)	0.506	BRFSS, NHIS (2018)
$\tilde{p}_e$	E-liquid price ($/ml)	3.90	E-cigarette User Survey
$\tilde{\tau}^c$	Average cigarette tax ($/pack)	2.92	Tax Policy Center (2019), ACS
$\tilde{\tau}^e$	Average e-liquid tax ($/ml)	0.233	Tax Foundation, RMS, Census

Notes: This table summarizes the parameters used for policy analysis. All dollar values are inflated to 2018 dollars. BRFSS, NHIS, MTF, and NYTS refer to sample surveys described in Table 1. Cigarette and e-liquid tax rates are averages across all U.S. states, weighted by population; the cigarette tax includes the federal cigarette tax of $1.01 per pack.

Table 5: **Optimal Tax and Welfare Effects of a Ban under Alternative Assumptions**

(a) **Optimal E-cigarette Tax ($/ml)**

Parameter assumptions	(1) $\alpha_{\text{adult}} = 0.29$ (p25)	(2) $\alpha_{\text{adult}} = 0.92$ (mean)	(3) $\alpha_{\text{adult}} = 1.64$ (p75)
1. Primary specification	6.41	18.26	31.91
2. $\sigma_\theta = \hat{\sigma}$ from Nielsen RMS	4.40	16.25	29.93
3. $\sigma_\theta =$ combined $\hat{\sigma}_\theta$	5.20	17.05	30.74
4. Perfect complements	17.06	28.91	42.57
5. Perfect substitutes	-5.86	5.99	19.66
6. $\tau^c = \varphi^c$	5.60	17.46	31.13
7. Rescale distortions so $\varphi^c = \tilde{\tau}^c$	0.91	2.85	5.08
8. $\gamma^c = 1.8$	1.68	5.01	8.84
9. $\gamma^c = 80$	30.42	85.61	149.21
10. $s_{\text{adult}} = 0$, $s_{\text{youth}} = 1$	10.22	31.09	55.12
11. $s_{\text{adult}} = 1$, $s_{\text{youth}} = 0$	5.58	15.47	26.86
12. σ_θ and η from Nielsen RMS without time trends, Viscusi (2016) φ^e	-55.82	-55.82	-55.82

(b) **Welfare Effects of E-cigarette Ban ($/person-year)**

Parameter assumptions	(1) $\alpha_{\text{adult}} = 0.29$ (p25)	(2) $\alpha_{\text{adult}} = 0.92$ (mean)	(3) $\alpha_{\text{adult}} = 1.64$ (p75)
1. Primary specification	24.84	90.53	166.26
2. $\sigma_\theta = \hat{\sigma}$ from Nielsen RMS	13.74	79.56	155.51
3. $\sigma_\theta =$ combined $\hat{\sigma}_\theta$	18.15	84.00	159.96
4. Perfect complements	84.17	149.99	225.89
5. Perfect substitutes	-43.16	22.61	98.45
6. $\tau^c = \varphi^c$	20.50	86.24	162.03
7. Rescale distortions so $\varphi^c = \tilde{\tau}^c$	-5.63	5.17	17.61
8. $\gamma^c = 1.8$	-1.24	17.21	38.49
9. $\gamma^c = 80$	158.32	464.79	818.25
10. $s_{\text{adult}} = 0$, $s_{\text{youth}} = 1$	90.88	319.88	583.71
11. $s_{\text{adult}} = 1$, $s_{\text{youth}} = 0$	18.29	67.77	124.82
12. σ_θ and η from Nielsen RMS without time trends, Viscusi (2016) φ^e	-323.02	-323.02	-323.02
13. $\eta = -.5$	12.38	78.15	153.95
14. $\eta = -1$	23.30	89.09	164.93
15. $\eta_{\text{youth}} = 1.5 \times \eta_{\text{adult}}$	27.61	93.31	169.09
16. Use $\tilde{p}_e$ from RMS	18.54	84.13	159.86
17. $\tau^e = \varphi^e$	-4.89	-4.89	-4.89
18. Double τ^e	23.62	87.49	161.48
19. Combined assumptions: $\eta = -0.5$, half distortions, and $\sigma = -\hat{\sigma}$	-9.43	23.46	61.37

Notes: Panel (a) presents estimates of the optimal e-cigarette tax using Equation (14). Panel (b) presents estimates of the welfare effects of an e-cigarette ban relative to current tax rates using Equation (15). The three columns present results at the 25th percentile, mean, and 75th percentile of α_{adult}, the ratio of uninternalized harms from vaping versus smoking. Each row varies a specific parameter assumption, and all other parameters are drawn from their distributions.

Figure 1: **National E-cigarette and Cigarette Sales over Time**

(a) **Combustible Cigarettes**

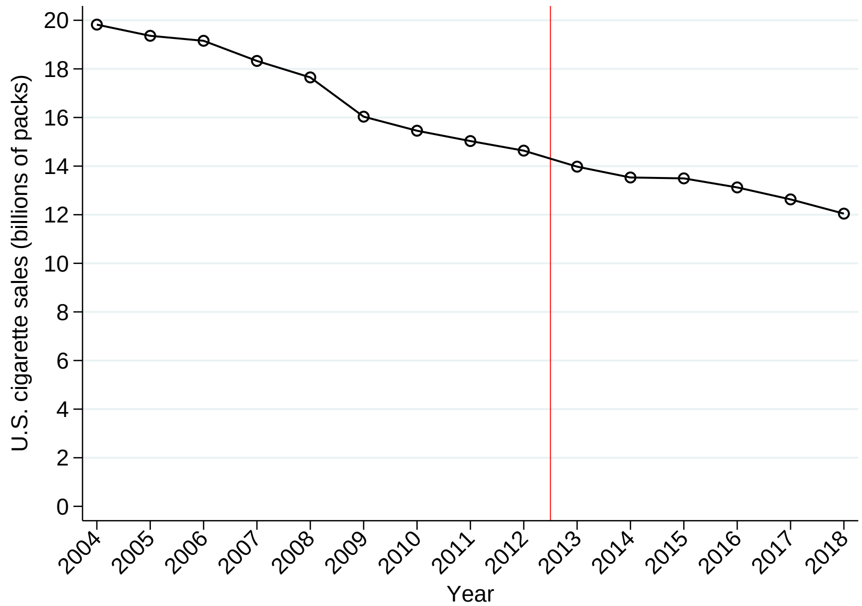

(b) **E-cigarettes**

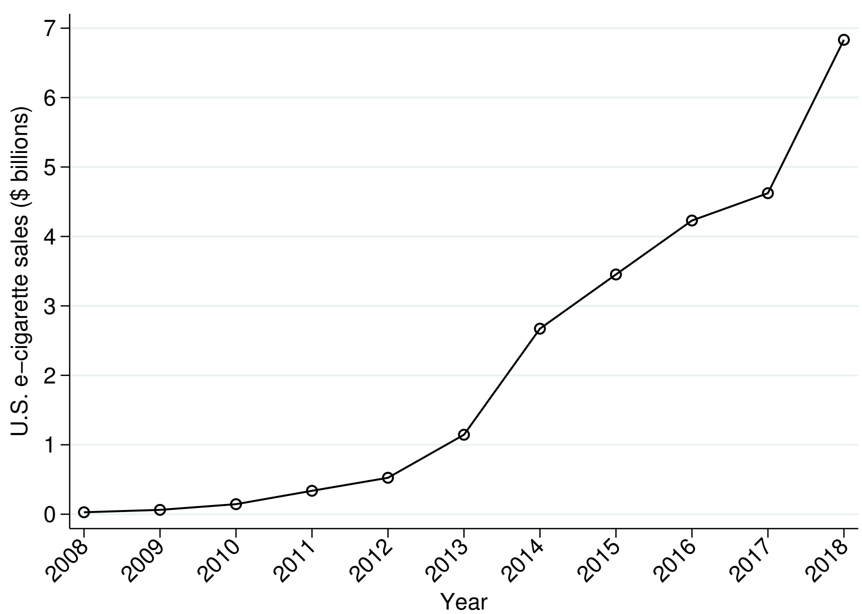

Notes: Data are from the Euromonitor Passport Cigarette and E-Vapour Products Databases.

Figure 2: **Smoking and Vaping Trends**

(a) **Adults**

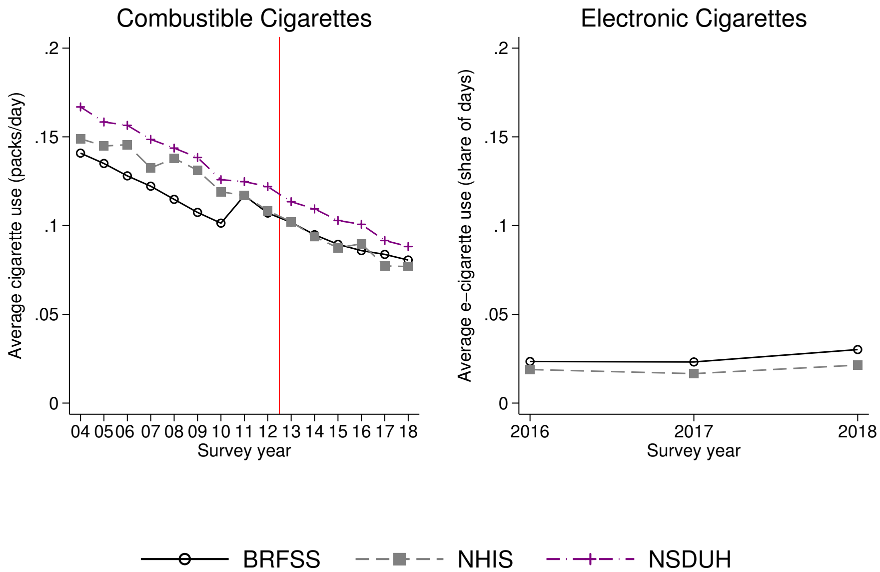

(b) **Youth**

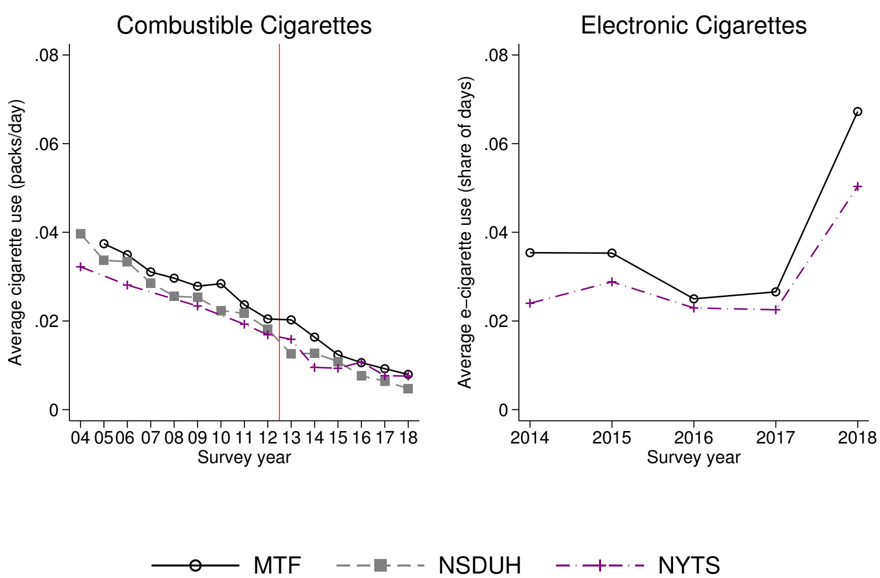

Notes: This figure presents combustible cigarette and e-cigarette use by survey and year. The BRFSS sampling frame changes in 2011, causing a jump in reported cigarette use. The NSDUH does not record data on vaping.

Figure 3: **Event Study of E-cigarette Tax Changes**

(a) **First Stage: Effect on ln(Price)**

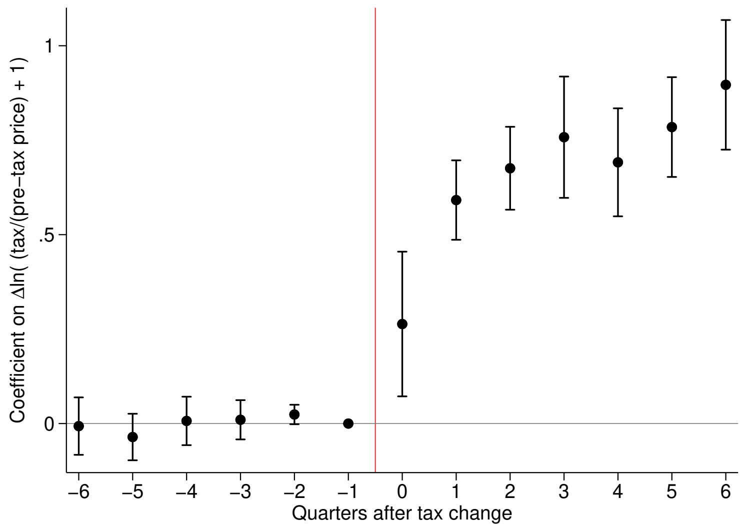

(b) **Reduced Form: Effect on ln(Quantity Sold + 1)**

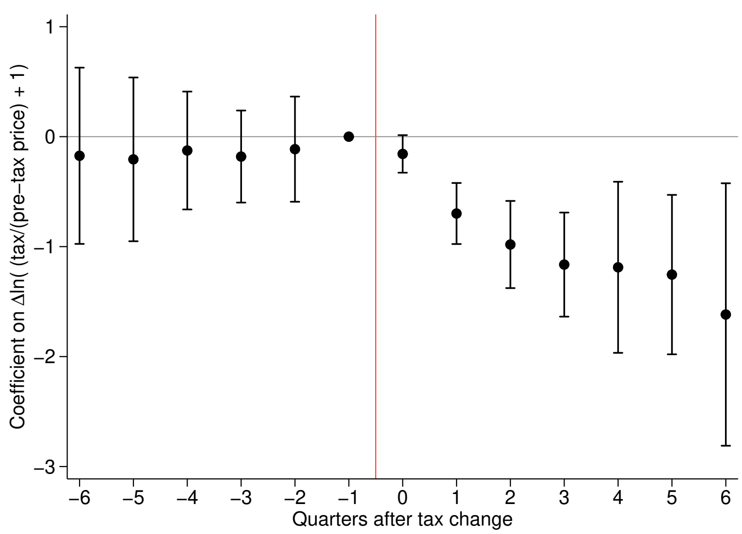

Notes: This figure presents estimates of the η_q parameters from Equation (17), an event study of the effects of e-cigarette tax changes. Panel (a) presents the first stage regression of ln(e-cigarette price) on the change in the log tax variable. Panel (b) presents the reduced form regression of the ln(e-cigarette units sold) on the change in the log tax variable.

Figure 4: **Demographic Predictors of Vaping**

(a) **Adults**

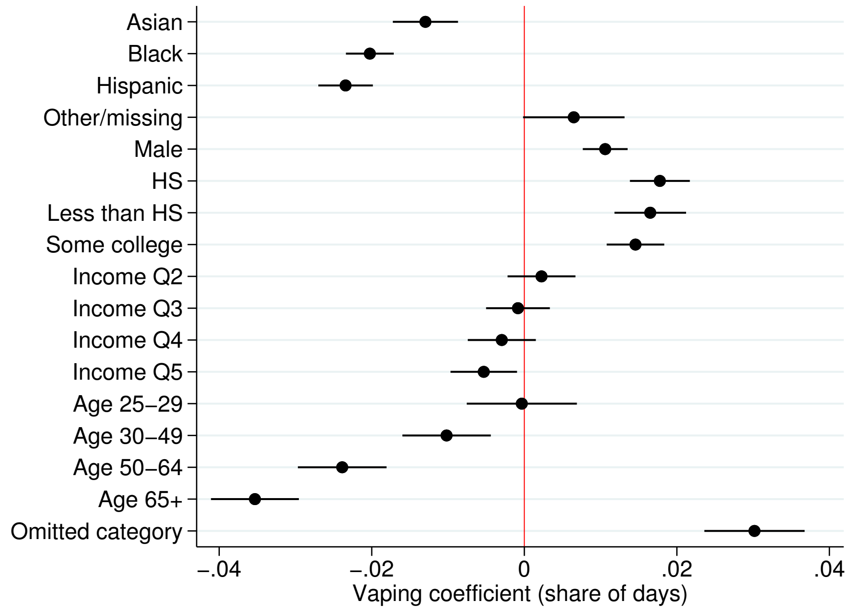

(b) **Youth**

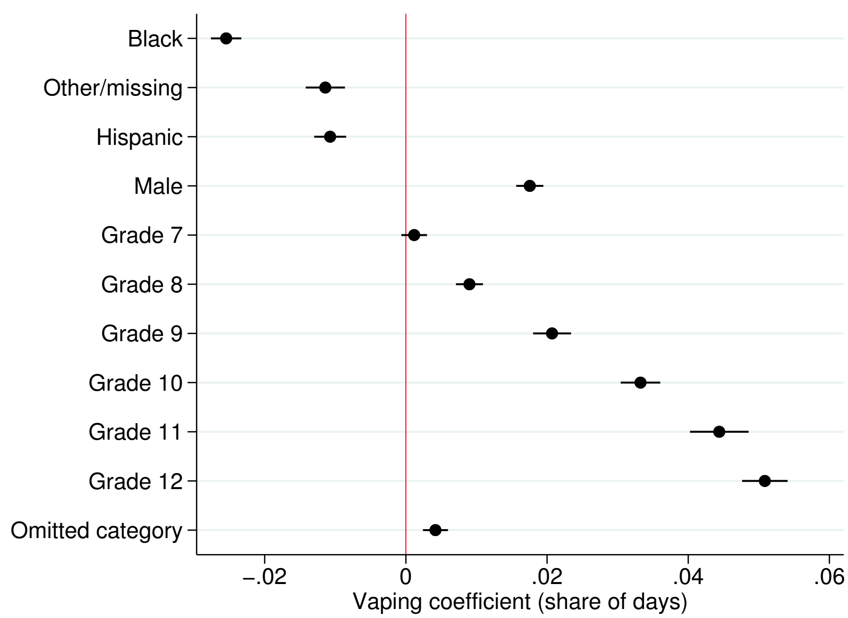

Notes: These figures present coefficients from Equation (20), a regression of vaping on demographic indicators. For adults, the omitted categories are White, female, college graduate, the lowest income quintile, and age group 18–24. For youth, the omitted categories are White, female, and grade 6. Panel (a) pools 2016–2018 data from BRFSS and NHIS; Panel (b) pools 2014–2018 data from MTF and NYTS. Standard errors are clustered by demographic cell.

Figure 5: **Smoking and Vaping Trends for High- versus Low-Vaping Demographics**

(a) **Adults**

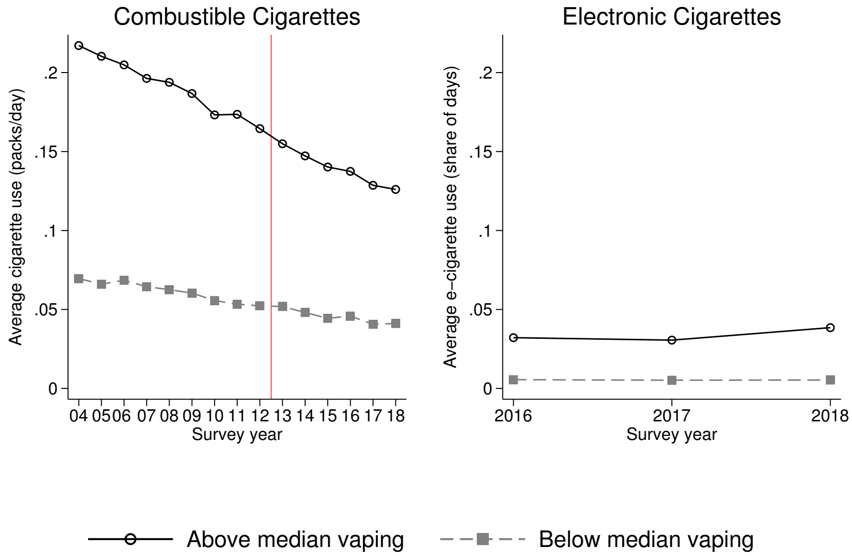

(b) **Youth**

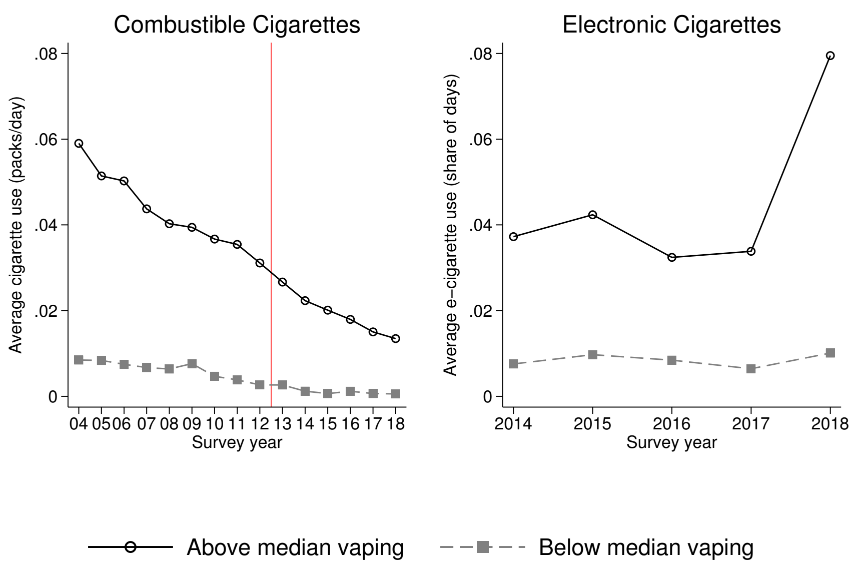

Notes: These figures present combustible cigarette and e-cigarette use for demographics with above- versus below-median predicted vaping, as predicted by Equation (20). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

Figure 6: **Difference in Smoking Trends for High versus Low Predicted Vaping**

(a) **Adults**

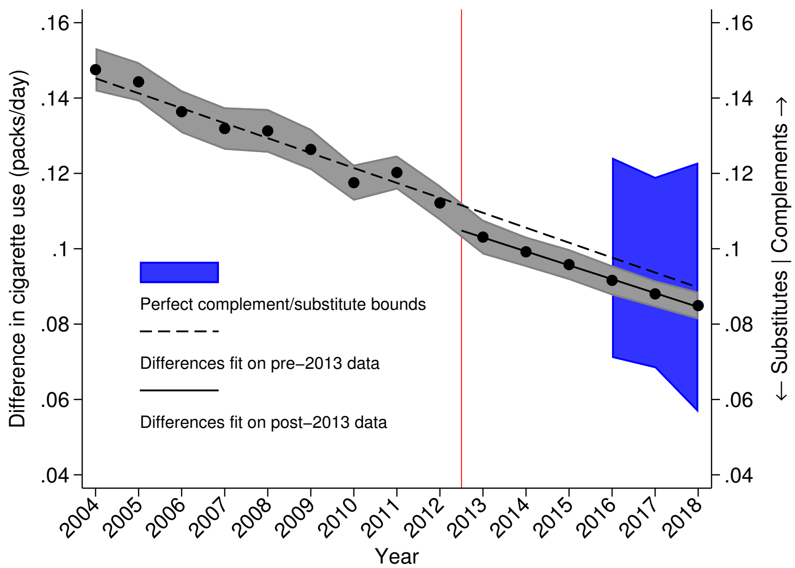

(b) **Youth**

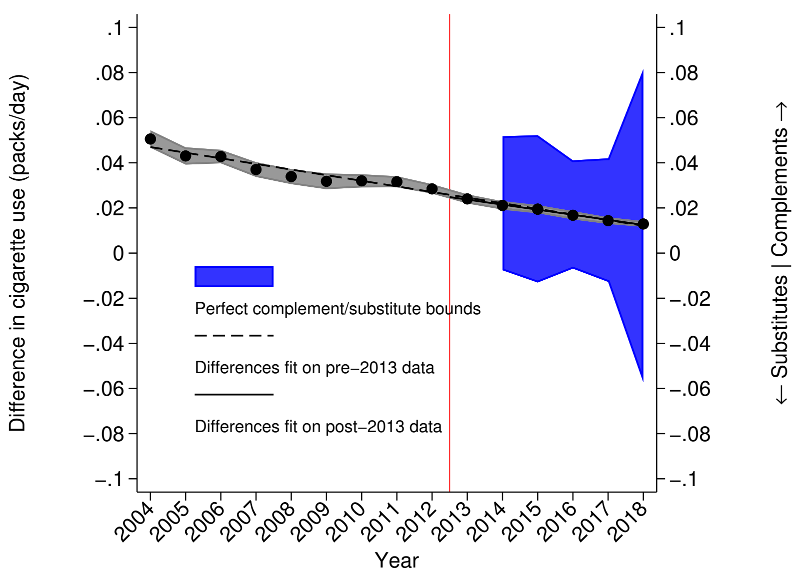

Notes: These figures present the difference in cigarette use for demographics with above- versus below-median predicted vaping, as predicted by Equation (20). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. The perfect complement (substitute) bounds show the difference in cigarette use that would be expected if daily vaping were a perfect complement (substitute) for smoking one pack of cigarettes per day. To construct perfect complement (substitute) bounds, we predict the difference in cigarette use using the pre-2013 time trend, then add (subtract) the difference in share of days vaped.

Figure 7: **Event Study of E-cigarette Introduction**

(a) **Adults**

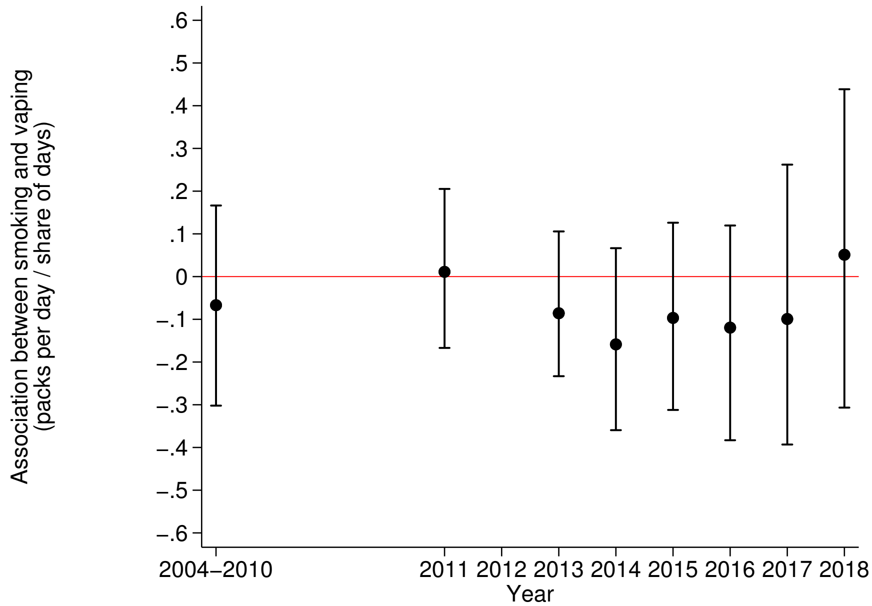

(b) **Youth**

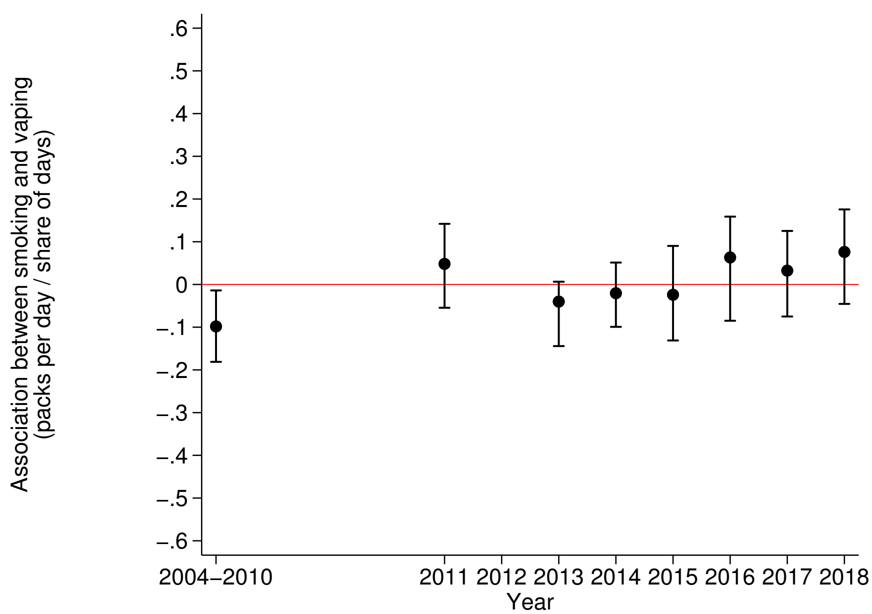

Notes: These figures present estimates of $\boldsymbol{\zeta}_t$ from Equation (24), a regression of cigarette use on predicted vaping interacted with year indicators, controlling for linear time trends and other controls. We estimate one indicator for the 2004–2010 period, and 2012 is the omitted year category. The confidence intervals reflect the 2.5th and 97.5th percentiles of estimates from 200 bootstrap replications, where we draw bootstrap samples by demographic cell.

Figure 8: **Substitution Parameters and Robustness Checks**

(a) **Adults**

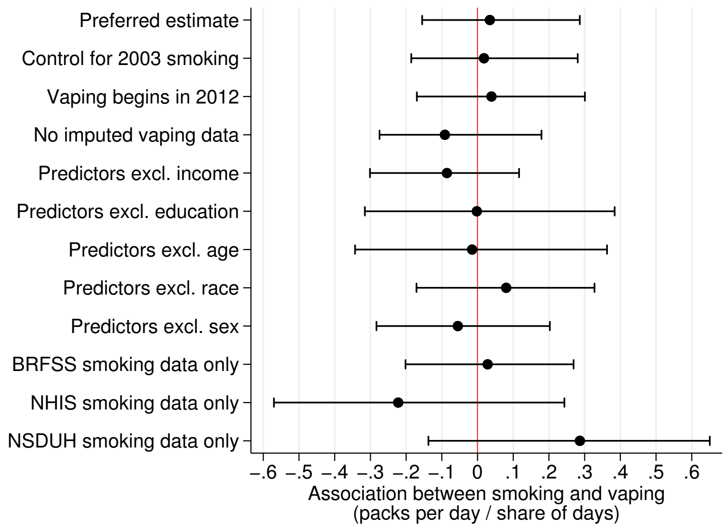

(b) **Youth**

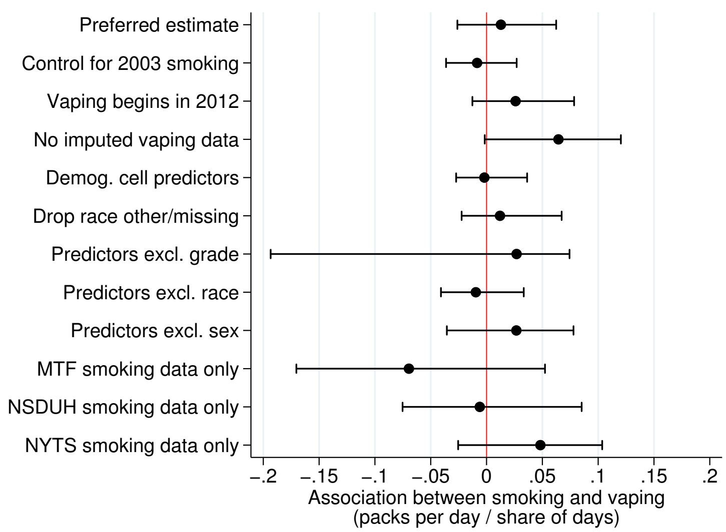

Notes: These figures present estimates of σ from Equation (21), a regession of smoking on predicted vaping controlling for controlling for linear time trends and other controls. *Control for 2003 smoking* includes additional controls for the 2003 cigarette use in person i's demographic cell and the interaction of that variable with a linear time trend. *Vaping begins in 2012* assumes zero vaping for all years before 2012 (instead of 2013 in the preferred estimate) and imputes vaping beginning in 2012 (instead of 2013). *Demog. cell predictors* uses demographic cells, rather than linear demographic groups, in $\boldsymbol{G}_i$. *Drop race other/missing* drops all observations with "other" or missing race/ethnicity. *No imputed vaping data* uses only observed vaping instead of imputing missing data beginning in 2013. *Predictors excl. age (or race, etc.)* omits age (or race, etc.) from the predictors in Equation (20). *BRFSS (or NHIS, etc.) smoking data only* uses only BRFSS (or NHIS, etc.) data when estimating Equation (5). *Drop race other/missing* drops all youth whose race/ethnicity is not Black, Hispanic, or White from both the predicted vaping and the smoking effects regressions. The confidence intervals reflect the 2.5th and 97.5th percentiles of estimates from 200 bootstrap replications, where we draw bootstrap samples by demographic cell.

Figure 9: **Expert Survey: Internalities and Externalities from Vaping Relative to Smoking**

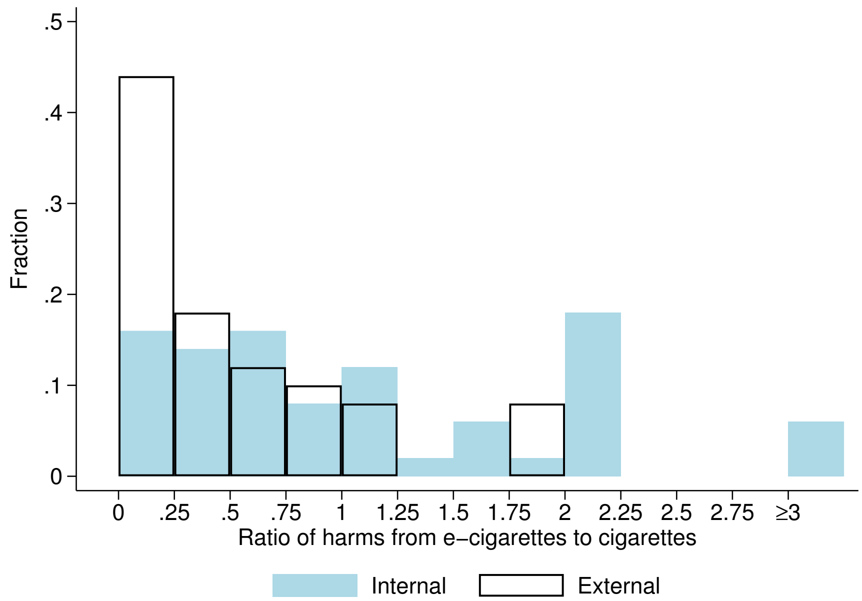

Notes: In our expert survey, we elicited the ratios of internalities and externalities from vaping relative to smoking. This figure presents the distributions of those ratios across experts.

Figure 10: **Optimal E-cigarette Tax: Distribution of Monte Carlo Simulations**

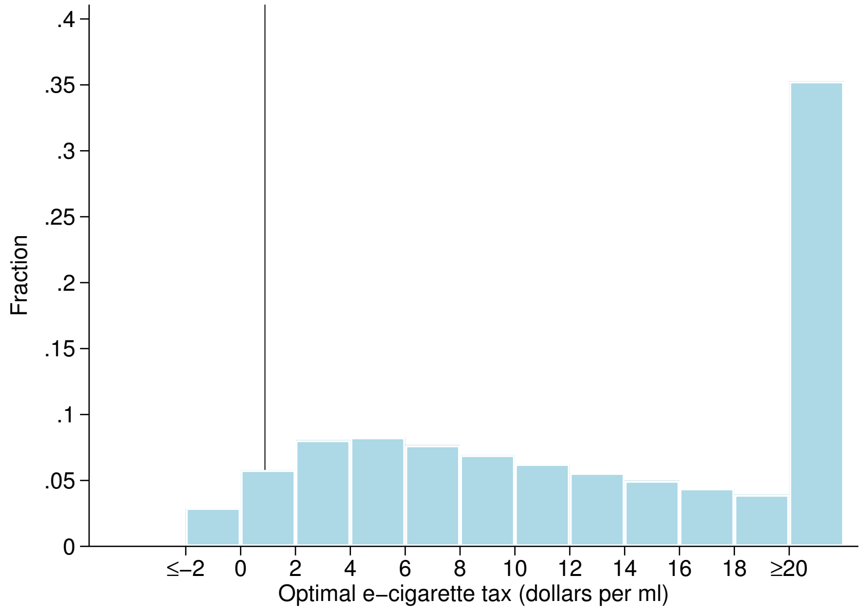

Notes: This figure presents the distribution of optimal e-cigarette taxes from Equation (14) over the distribution of Monte Carlo simulations. The vertical line at \$0.89/ml represents the average existing e-cigarette tax rate.

Figure 11: **Welfare Effects of E-cigarette Ban: Distribution of Monte Carlo Simulations**

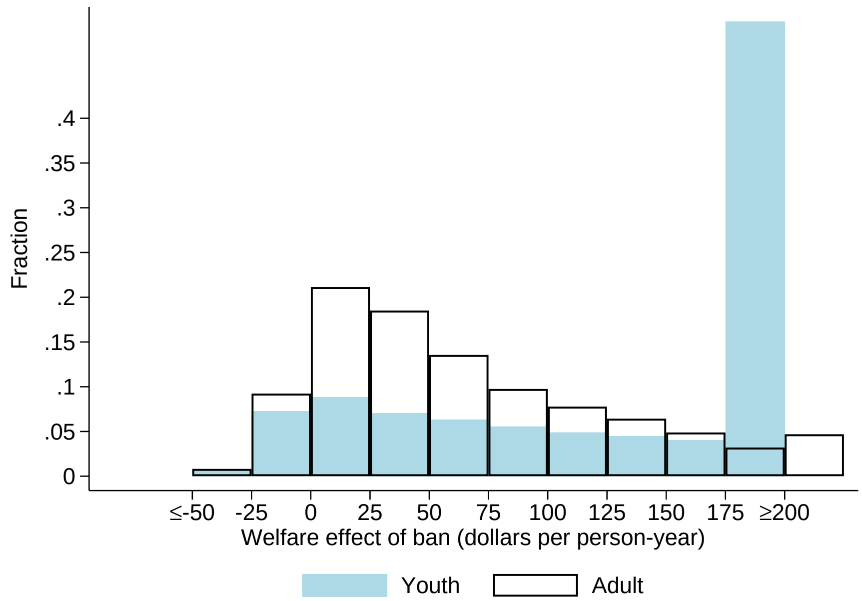

Notes: This figure presents the welfare effects of an e-cigarette ban compared to current tax rates from Equation (15) over the distribution of Monte Carlo simulations. We present separate estimates for youth-specific and adult-specific bans.

Figure 12: **Contribution of Parameters to Policy Uncertainty**

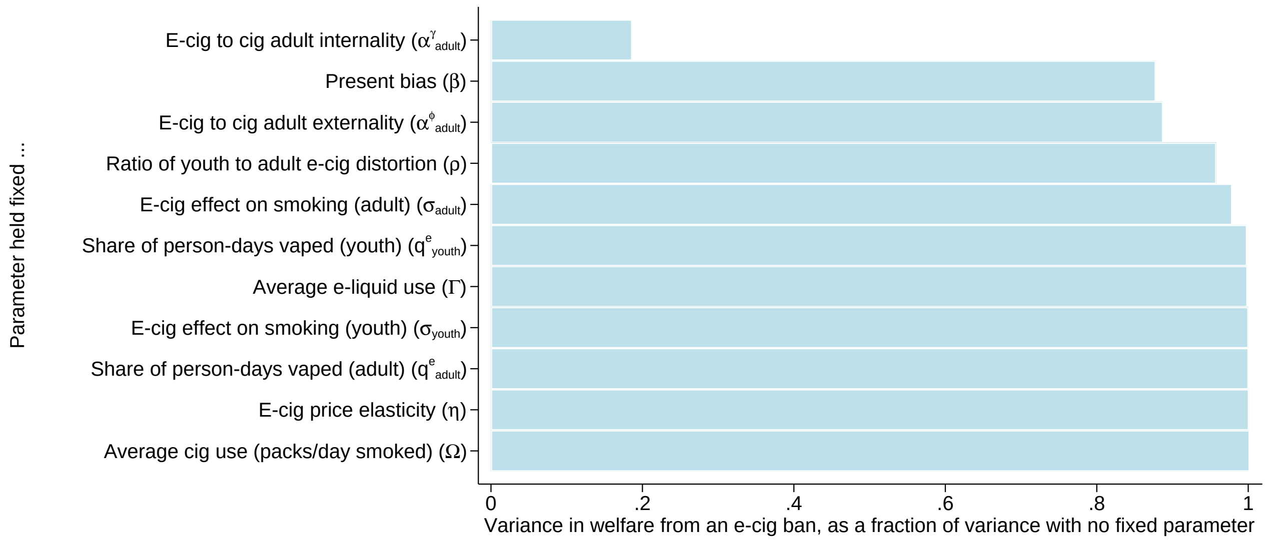

Notes: This figure presents the variance across Monte Carlo simulations of the welfare effects of an e-cigarette ban from Equation (15), holding the reported parameter fixed at its mean.

Figure 13: **Welfare Effects of E-cigarette Ban**

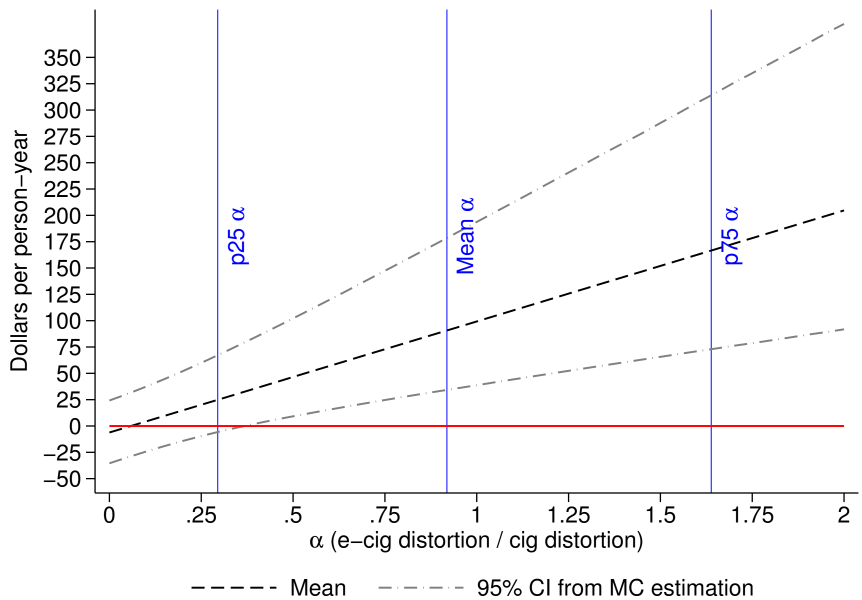

Notes: This figure presents the mean and 95 percent confidence interval of the welfare effects of an e-cigarette ban from Equation (15) over the distribution of Monte Carlo simulations, for different values of α_a, the ratio of uninternalized harms from daily smoking versus daily vaping. The bounds are not symmetric because the e-cigarette own-price elasticity η enters Equation (15) in the denominator.

Appendix

For Online Publication

Optimal Regulation of E-cigarettes: Theory and Evidence

Hunt Allcott and Charlie Rafkin

A Theory Appendix

A.1 Optimal Taxes

After substituting the utility function and consumer budget constraint, social welfare at time 0 is

$$W(\boldsymbol{\tau}) = \sum_{\theta,t} \delta^t s_\theta \left[u_\theta\left(\boldsymbol{q}_{\theta t}; S_t\right) - \boldsymbol{p} \cdot \boldsymbol{q}_{\theta t} + z_{\theta t} + T_t\right]. \tag{25}$$

Substituting in the balanced budget constraint $T_t = \sum_\theta \left(\boldsymbol{\tau} - \boldsymbol{\phi}_\theta\right) \cdot \boldsymbol{q}_{\theta t}$ gives

$$W(\boldsymbol{\tau}) = \sum_{\theta,t} \delta^t s_\theta \left[u_\theta\left(\boldsymbol{q}_t; S_t\right) - \boldsymbol{p} \cdot \boldsymbol{q}_t + z_{\theta t} + \left(\boldsymbol{\tau} - \boldsymbol{\phi}_\theta\right) \cdot \boldsymbol{q}_{\theta t}\right]. \tag{26}$$

The effect of a marginal change in q_t^k on type θ's value function is the effect on current period utility, $\frac{\partial u_\theta(\boldsymbol{q}_{\theta t};S_t)}{\partial q_t^k} - p^k$, plus the discounted effect on the continuation value, $\delta \frac{\partial V_\theta(S_{t+1})}{\partial S_{t+1}} \cdot \frac{\partial S_{t+1}}{\partial q_t^k}$. Thus, recalling that $\boldsymbol{p}$ is the tax-inclusive price, the derivative of social welfare with respect to τ^j is

$$\begin{aligned}
\frac{\partial W_r(\boldsymbol{\tau})}{\partial \tau^j} &= \sum_{\theta,t,k} \delta^t s_\theta \left[\left(\frac{\partial u_\theta\left(\boldsymbol{q}_{\theta t}; S_t\right)}{\partial q_t^k} + \delta \frac{\partial V_\theta\left(S_{t+1}\right)}{\partial S_{t+1}} \cdot \frac{\partial S_{t+1}}{\partial q_t^k} - p^k\right) \frac{dq_t^k}{d\tau^j} - q_{\theta t}^k + \left(\tau^k - \phi_\theta^k\right) \frac{dq_{\theta t}^k}{d\tau^j} + q_{\theta t}^k\right] \\
&= \sum_{\theta,t,k} \delta^t s_\theta \left[-\gamma_\theta^k(\boldsymbol{p}, S_t) \frac{dq_{\theta t}^k}{d\tau^j} + \left(\tau^k - \phi_\theta^k\right) \frac{dq_{\theta t}^k}{d\tau^j}\right] \\
&= \sum_{\theta,t,k} \delta^t s_\theta \left(\tau^k - \varphi_\theta^k(\boldsymbol{p}, S_t)\right) \frac{dq_{\theta t}^k}{d\tau^j},
\end{aligned} \tag{27}$$

where the second line follows from the definition of $\gamma_\theta^j(\boldsymbol{p}, S_t)$ in Equation (5) and the third line follows from the definition of $\varphi_\theta^k(\boldsymbol{p}, S_t)$ in Equation (9). Setting equal to zero and re-arranging gives

$$\tau^j \sum_{\theta,t} \delta^t s_\theta \frac{dq_{\theta t}^j}{d\tau^j} = \sum_{\theta,t} \delta^t s_\theta \varphi_\theta^j(\boldsymbol{p}, S_t) \frac{dq_{\theta t}^j}{d\tau^j} + \sum_{\theta,t} \delta^t s_\theta \left(\varphi_\theta^{-j}(\boldsymbol{p}, S_t) - \tau^{-j}\right) \frac{dq_{\theta t}^{-j}}{d\tau^j}, \tag{28}$$

and dividing by $\sum_{\theta,t} \delta^t s_\theta \frac{dq_{\theta t}^j}{d\tau^j}$ gives Equation (10).

A.2 Welfare Effect of a Ban

The welfare effect of banning e-cigarettes beginning in period 0 is

$$\begin{aligned}
\Delta W &= \int_{\tilde{\tau}^e}^{\infty} \frac{\partial W(\tau)}{\partial \tau^e} d\tau^e \\
&= \int_{\tilde{\tau}^e}^{\infty} \sum_{\theta,t,j} \delta^t s_\theta \left(\tau^j - \varphi_\theta^j(\boldsymbol{p}, S_t)\right) \frac{dq_{\theta t}^j}{d\tau^e} d\tau^e \\
&= \sum_{\theta,t,j} \delta^t s_\theta \left[\int_{\tilde{\tau}^e}^{\infty} \tau^j \frac{dq_{\theta t}^j}{d\tau^e} d\tau^e - \int_{\tilde{\tau}^e}^{\infty} \varphi_\theta^j(\boldsymbol{p}, S_t) \frac{dq_{\theta t}^j}{d\tau^e} d\tau^e \right]. \qquad (29)
\end{aligned}$$

Integrating by parts gives

$$\sum_j \int_{\tilde{\tau}^e}^{\infty} \tau^j \frac{dq_{\theta t}^j}{d\tau^e} d\tau^e = \sum_j \tau^j q_{\theta t}^j \bigg|_{\tilde{\tau}^e}^{\infty} - \int_{\tilde{\tau}^e}^{\infty} q_{\theta t}^e d\tau^e = \sum_j \tilde{\tau}^j \Delta q_{\theta t}^j - \int_{\tilde{\tau}^e}^{\infty} q_{\theta t}^e d\tau^e. \qquad (30)$$

Substituting Equations (12) and (30) into Equation (29) gives

$$\Delta W = \sum_{\theta,t} \delta^t s_\theta \left[- \int_{\tilde{\tau}^e}^{\infty} q_{\theta t}^e d\tau^e + \sum_j \tilde{\tau}^j \Delta q_{\theta t}^j - \sum_j \overline{\varphi}_{\theta t}^j \Delta q_{\theta t}^j \right].$$

Re-arranging gives Equation (13).

A.3 Empirical Implementation

We impose two assumptions to estimate both the optimal tax and the welfare effect of a ban.

Assumption 1. Homogeneous and constant own-price elasticity: $\eta_{\theta t}^j = \eta^j$, *for all* (θ, t).

Assumption 2. Zero covariance: $\varphi_\theta^j(\boldsymbol{p}, S_t)$, $\sigma_{\theta t}^j$, $q_{\theta t}^j$, *and* t *have pairwise zero covariance conditional on* θ.

Optimal tax. Define $\eta^j = \frac{dq_{\theta t}^j/dp^j}{q_{\theta t}^j/p^j}$ and $\sigma_{\theta t}^j := \frac{dq_{\theta t}^{-j}/dp^j}{dq_{\theta t}^j/dp^j}$ as the own-price elasticity and substitution parameters. The η and $\sigma_{\theta t}$ defined in Section 1 are for $j = e$. Since $\eta^j = \frac{dq_{\theta t}^j/dp^j}{q_{\theta t}^j/p^j}$, we have $\frac{dq_{\theta t}^j}{dp^j} = \eta^j q_{\theta t}^j/p^j$ and $\frac{dq_{\theta t}^{-j}}{dp^j} = \sigma_{\theta t}^j \eta^j q_\theta^j/p^j$. Under Assumption 1, the optimal tax from Equation (10) becomes

$$\tau^{*j} = \frac{\sum_{\theta,t} \delta^t s_\theta q_{\theta t}^j \varphi_\theta^j(\boldsymbol{p}, S_t)}{\sum_{\theta,t} \delta^t s_\theta q_{\theta t}^j} + \frac{\sum_{\theta,t} \delta^t s_\theta q_{\theta t}^j \sigma_{\theta t}^j \left(\varphi_\theta^{-j}(\boldsymbol{p}, S_t) - \tau_t^{-j}\right)}{\sum_{\theta,t} \delta^t s_\theta q_{\theta t}^j}. \qquad (31)$$

Adding Assumption 2 yields

$$\tau^{*j} = \frac{\sum_{\theta} s_\theta q_\theta^j \varphi_\theta^j}{\sum_{\theta} s_\theta q_\theta^j} + \frac{\sum_{\theta} s_\theta q_\theta^j \sigma_\theta^j \left(\varphi_\theta^{-j} - \tau_t^{-j}\right)}{\sum_{\theta} s_\theta q_\theta^j}. \tag{32}$$

Welfare effect of ban. We add a further functional form assumption to identify the perceived consumer surplus change.

Assumption 3. Perceived consumer surplus change: $-\int_{\tilde{\tau}^e}^{\infty} q_\theta^e(\boldsymbol{p}) d\tau^e = \Delta q_\theta^e \frac{\tilde{p}^e}{-2\eta}$.

Under Assumption 3, Equation (13) becomes

$$\Delta W = \sum_{\theta,t} \delta^t s_\theta \left[\Delta q_{\theta t}^e \frac{\tilde{p}^e}{-2\eta} - \sum_j \Delta q_{\theta t}^j \left(\overline{\varphi}_\theta^j(\boldsymbol{p}, S_t) - \tau^j\right)\right]. \tag{33}$$

Adding Assumption 2 yields

$$\Delta W = \frac{1}{1-\delta} \sum_{\theta} s_\theta \left[\Delta q_\theta^e \frac{\tilde{p}^e}{-2\eta} - \sum_j \Delta q_\theta^j \left(\varphi_\theta^j - \tau^j\right)\right]. \tag{34}$$

Multiplying by $1-\delta$ gives the average per-period welfare effect:

$$\Delta \overline{W} = \sum_{\theta} s_\theta \left[\Delta q_\theta^e \frac{\tilde{p}^e}{-2\eta} - \sum_j \Delta q_\theta^j \left(\varphi_\theta^j - \tau^j\right)\right]. \tag{35}$$

B Data Appendix

B.1 RMS Data

B.1.1 Data Construction

We construct two datasets: (1) a UPC-cluster-month dataset of *e-cigarette* units sold and prices data, and (2) a UPC-cluster-month dataset of *cigarette* units sold and prices data.

Sample restrictions. We exclude data from stores that are not observed for the full 2013–2017 sample period. Since UPCs with low sales are more likely to enter and exit the sample and create an unbalanced panel, we drop UPCs with less than $100,000 in total sales from the analysis sample.

Weeks that occur in two months are assigned to the later month (i.e., the month in which the week's Saturday falls).

Weights. For simplicity, we refer to our estimates as being weighted by sales, but we do not weight by raw sales because sales are endogenous to the tax rate. We construct e-cigarette weights as follows. We construct the total sales for a given UPC-year that occur in states without e-cigarette taxes. We then divide this number by the total e-cigarette sales that occur in untaxed states in that year. Cigarette sales are nearly always subject to some tax. To construct weights for cigarette analyses, we construct the total sales in a given UPC-year (excluding that observation's own UPC-year-cluster sales), as a fraction of the total sales in that year across UPCs (excluding sales in the given UPC-year-cluster). We exclude the observation's own UPC-cluster-year sales from the numerator and denominator to account for the fact that sales are endogenous to the tax environment.

E-cigarette dataset. We construct unit-weighted prices at the UPC-cluster-month level. The cigarette prices in this dataset are cluster-month unit-weighted cigarette post-tax prices, including the monthly cigarette sales tax per pack. The cigarette tax rate is the state and national cigarette tax in a given state-month, divided by the unit-weighted cigarette post-tax price less the state-month cigarette tax.

Cigarette dataset. We convert Nielsen units and prices per unit to packs. We construct unit-weighted prices at the UPC-cluster-month level. The cigarette tax rate is the state and national cigarette tax as a fraction of the observation's unit-weighted UPC-month cigarette post-tax price less the state cigarette tax, excluding the UPC's own cluster. We drop observations where the official cigarette tax is more than the scanner post-tax price. We construct unit-weighted cluster-month e-cigarette prices, and we obtain the e-cigarette tax by using the algorithm in the following subsection. Since we are working with cluster-month data, we use the sales-weighted e-cigarette size across all clusters and the unit-weighted price across untreated clusters.

B.1.2 Constructing the E-cigarette Tax Variable

There are two types of e-cigarette taxes: ad-valorem taxes (where the tax is a percentage of the UPC price) and specific taxes (where the tax is a constant per milliliter of e-liquid). In all clusters, taxes collected are included in the UPC price recorded in RMS. Let τ'_{st} represent the ad-valorem tax rate in cluster s. With full pass-through, $\tau_{kst} = \tau'_{st}$ in ad-valorem cluster-months, for all UPCs k. To construct a consistent instrument that appropriately scales the magnitude of the tax across different regimes, we convert specific taxes to ad-valorem taxes. For each UPC-month, we generate the unit-weighted price p'_k, across all months, using only clusters with no e-cigarette taxes. Let size_k denote the mililiters of e-liquid contained in UPC k. The ad-valorem tax for UPC k in a cluster s with a specific tax α_{st} per mililiter of e-liquid in month t is given by $\tau_{kst} = \frac{\alpha_{st} \cdot \text{size}_k}{p'_k}$. In the final analysis, we drop the 0.12% of the total observations have $\tau_{kst} > 1$ or for which we do not observe any sales in states with no e-cigarette taxes (to construct p'_k). Summarizing,

$$\tau_{kst} = \left\{ \begin{array}{cc} 0, & s \text{ has no e-cigarette tax} \\ \tau'_{st}, & s \text{ has an ad-valorem e-cigarette tax} \\ \frac{\alpha_{st} \cdot \text{size}_k}{p'_k}, & s \text{ has a specific e-cigarette tax} \end{array} \right\}.$$

The RMS data do not consistently record the size, in milliliters of liquid, of vaping products. We begin with the list of UPC sizes generously shared by the authors of Cotti et al. (2020). We augment their list with hand-collected information on the mililiters of liquid for the largest UPCs. For UPCs where we could accurately record size, we convert the per-ml taxes to taxes that are a fraction of the UPC price. In the final dataset, we observe 79 percent of the observations' sizes. For other UPCs, we convert prices to the average sales-weighted size for UPCs whose size we did record.

The city of Chicago enacted a separate tax several months before Cook County. Because we only observe the county in which sales take place, we assume that: (i) taxes that occur in Chicago apply throughout Cook County, Illinois, and: (ii) the Cook County tax was additive on top of the Chicago tax. Moreover, Chicago enacted a tax of \$0.80 per unit or \$0.55 per ml of e-liquid. Because of the difficulty in converting RMS units to the units taxed, we assume Chicago's tax is per ml of e-liquid.

In the event study analysis, we construct a variable τ'_{kstq} that varies by UPC, cluster, calendar month, and event quarter. In months prior to treatment in specific tax states, where τ_{ksq} varies by $k$and q, we construct α_{s0}, the size of the specific tax in cluster s in event-month 0, and generate $\tau_{kstq} = \frac{\alpha_{s0} \cdot \text{size}_k}{p'_k}$.[24]

Table A1: **E-cigarette Tax Changes Through 2017**

Area (state, county, or city)	Date	Tax rate
California	4/2017, 7/2017	27.3%, 65% of wholesale price
Chicago, IL	1/2016	\$0.80 per unit / \$0.55 per ml
Cook County, IL	5/2016	\$0.20 per ml
Kansas	7/2016, 7/2017	\$0.20, \$0.05 per ml
Louisiana	7/2015	\$0.05 per ml
Minnesota	8/2010, 7/2013	35%, 95% of wholesale price
Montgomery County, MD	8/2015	30% of wholesale price
North Carolina	6/2015	\$0.05 per ml
Pennsylvania	7/2016	40% of retail price
Washington, DC	10/2015, 10/2016	67%, 65% of wholesale price
West Virginia	7/2016	\$0.075 per ml

Notes: Data are from Pesko, Courtemanche and Maclean (2019, Appendix Table 2) and Tax Foundation (2019). The table excludes changes in Alaska, which does not appear in the RMS data.

[24]For consistency with other sample restrictions, we drop the pre-treatment observations where the implied $\tau_{ksq} > 1$.

B.2 Sample Surveys

This section details our construction of harmonized samples across the BRFSS, MTF, NHIS, NSDUH, and NYTS.

B.2.1 Sample Weights

All surveys excluding MTF come with nationally representative sample weights; MTF provides relative sampling odds, which we transform to sample weights. We use the survey-provided sample weights for adults. For youth, we rescale the sampling weights by the sum of weights within dataset-grade-year grade. Hence, within dataset, each observation retains its sampling weight relative to other observations within the dataset. Once we append the datasets, the sampling weights are appropriately scaled with respect to one another.

B.2.2 Income quintile construction

We construct income quintile within dataset-year, including sampling weights. Income is often recorded in bins, and occasionally the bins cut across quintile cut points. We assign to the lower quintile except in the case of the NHIS's first quintile, because doing so would only four quintiles in some years. To ensure there are five income quintiles in every year, we re-assign incomes that cut across the first and second quintiles to income quintile 1 in the NHIS prior to 2006 and income quintile 2 for 2007–2018. In the 2018 NSDUH, there are only four income groups recorded, which we code as quintiles 1, 2, 4, and 5.

B.2.3 Adult Smoking (NHIS, NSDUH, BRFSS)

NHIS. We use the *smknow*, *cigsda1*, and *cigsda2* variables to identify people who report smoking "every day," "some days," or "not at all." Among people who smoke every day, we use *cigsda1* to construct the average number of cigarettes smoked per day. If someone reports smoking "not at all," we impose that these people smoke 0 cigarettes per day on all days. Among people who report smoking "some days," we use *cigdamo* to generate the average number of days smoked in the past 30 days and the *cigsda2* variable to generate the average number of cigarettes smoked on days when the person smokes; we extract the average number of cigarettes smoked per day as ci$gsda2 \times cigdamo/30$.

NSDUH. We use the *cig30av* variable to compute the average number of cigarettes smoked per day on days smoked. Because the variable is interval censored, we use the midpoint of the reported ranges. We code the final interval ("35 cigarettes or more, about two packs") as 50 cigarettes (2.5 packs), for consistency with other top-coded datasets. We use the *cig30use* variable to compute the average number of days in the past 30 days when the respondent smoked. Among the small proportion of people who do not remember the precise number of days smoked, we use the midpoint of ranges reported in the *cg30est* variable to compute an estimate of the number of days smoked. We extract the number of cigarettes smoked per day in the past 30 as (number of days smoked in the past 30 / 30) $\times$ (number of cigarettes smoked on days smoked).

BRFSS. We use the *smokeday* and *smokday2* variables to construct a variable encoding whether someone smokes "every day," "some days," or "not at all." We rescale these variables for comparability by using the following algorithm.

For each year in 2004-2018, append the NHIS and NSDUH datasets. Keep only those people with non-missing values of demographic variables used in the main analysis. Extract smoking intensity among "every day" smokers: compute the average number of cigarettes smoked per day among people who report smoking 30 days in the past 30 in the NSDUH, or who smoke "every day" in the NHIS. Extract smoking intensity among "sometimes" smokers: compute the average number of cigarettes smoked per day among people who report smoking between 1 and 29 days in the past 30 in the NSDUH or who smoke "some days" in the NHIS. Construct a "predicted" smoking intensity for that year and smoking status by regressing the number of cigarettes smoked on survey year (i.e, compute a linear fit). Weight regression by sampling weights in each dataset. Divide the number of cigarettes smoked by 20 to obtain number of packs consumed per day.

Among people who report smoking "every day" in BRFSS, we impose that the person smokes the average number of packs in that year among every day smokers. Among people who report smoking "some days" in BRFSS, we impose that the person smokes the average number of packs in that year among "sometimes" smokers.

B.2.4 Adult Vaping (NHIS, BRFSS)

NHIS. We use the *ecig30d2*, *ecigcur2*, and *ecigev2* variables to construct a variable that is 1 if the person vaped "every day" (in *ecigcur2*), 0 if the person vaped "not at all" (in *ecigcur2*) and is *ecig30d2*/30 if the person reports vaping "some days" (in *ecigcur2*).

BRFSS. We use the *ecignow* and *ecigaret* variables to construct a variable that encodes whether the person vapes "every day," "some days," or "not at all." We use a similar algorithm as for vaping to rescale the variable for comparability: Among people who report vaping "not at all" in BRFSS, impose that the person has a vaping equivalent of 0. Among people who report vaping "every day" in BRFSS, impose that the person has a vaping equivalent of 1. For each year in 2016–2018, append the NHIS datasets. Keep only those people with non-missing values of demographic variables used in the main analysis. Extract smoking intensity among "sometimes" vapers: compute the average number of days vaped in the past 30 among people who report vaping "some days" in the NHIS. Among people who report smoking "some days" in BRFSS, impose that the person has a vaping equivalent of the average value extracted among vapers who report vaping "some days." Unlike in the exercise for smoking, do not generate separate values for each year.

B.2.5 Youth Smoking (MTF, NYTS, NSDUH)

MTF. We define packs per day as the number of cigarettes smoked per day on average, divided by 20. We recode the top-coded observations that report smoking 2 or more packs per day as smoking 50 cigarettes per day.

NYTS. We use the midpoint of the interval containing the number of cigarettes per day smoked and the midpoint of the number of days smoked to obtain the number of packs smoked per day. We code "20 or more" cigarettes per day as 30 cigarettes per day.

NSDUH. Same as adults.

B.2.6 Youth Vaping (MTF, NYTS)

Both datasets. We extract the midpoint of the interval containing the number of times the respondent reports vaping electronic cigarettes last month. We define vaping equivalents as the midpoint of this interval, divided by 30.

Additional details about the MTF vaping data. The MTF has several different variables from 2014–2018 that record the number of days the respondent reports vaping. By year, they are as follows (emphasis from MTF codebooks).

2014:

- During the LAST 30 DAYS, on how many occasions (if any) have you used electronic cigarettes (e-cigarettes)?

2015:

- During the LAST 30 DAYS, on how many occasions (if any) have you used electronic cigarettes (e-cigarettes)?
- During the LAST 30 DAYS, on how many days (if any) have you used an electronic vaporizer such as an e-cigarette?

2016:

- During the LAST 30 DAYS, on how many days (if any) have you used electronic cigarettes (e-cigarettes)?
- During the LAST 30 DAYS, on how many days (if any) have you used an electronic vaporizer such as an e-cigarette?

2017:

- On how many occasions (if any) have you vaped NICOTINE during the last 30 days?

- During the LAST 30 DAYS, on how many occasions (if any) have you used electronic cigarettes (e-cigarettes)?

2018:

- On how many occasions (if any) have you vaped NICOTINE during the last 30 days?
- During the LAST 30 DAYS, on how many occasions (if any) have you used electronic cigarettes (e-cigarettes)?

We combine these reports as follows. If a respondent is ever recorded asked *multiple* vaping questions, we take the average. If the respondent records vaping more than 30 times in the past month, we recode this as 30 (such that the maximum number of *days* in the last month is 30). Figure A1 illustrates that mean vaping rates align well across these reports.

Figure A1: **MTF Vaping Rates by Question**

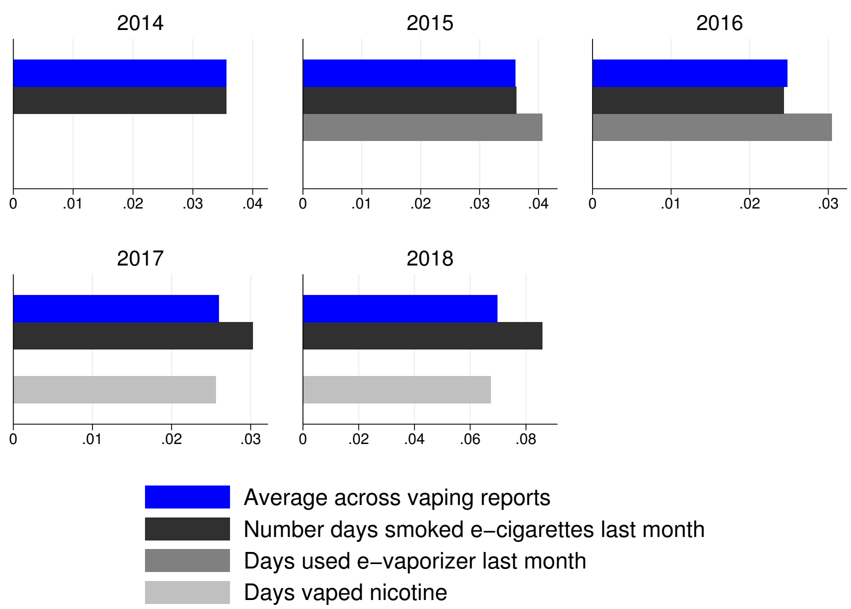

Notes: This figure presents vaping rates by year and question from the Monitoring the Future survey.

B.2.7 Additional Issues in Sample Surveys

NSDUH. The NSDUH is the sole youth survey that does not have a clean way of identifying students' current grade to provide comparability with MTF and NYTS. We therefore count people in grades 6–12, or people who are age 18, as youth. Because we include 18–24 year olds in the adult estimations, this means the 18 year-olds in the NSDUH appear in both the youth and adult surveys. The public-use NSDUH data also provide ages in bins that are not comparable to the

BRFSS and NHIS for some adults. For demographic controls, we code NSDUH 18–23 year olds as 18–24 year olds and NSDUH 24–29 year olds as 25–29 year olds.

BRFSS. Because of inconsistent data collection, we drop survey respondents from Guam, Puerto Rico, and other territories from the BRFSS sample.

MTF. The MTF samples only the 48 contiguous states. The MTF does not sample dropouts. We are limited to four race/ethnicity groups in the youth dataset because Asian is not a separate category from other race in the public-use MTF.

NYTS. The NYTS does not sample dropouts.

B.2.8 Total Quantities in Sample Surveys versus Sales Data

The total cigarette and e-cigarette sales implied by our sample survey data and unit conversion parameters line up reasonably closely with national sales data. Multiplying 2018 average smoking for adults and youths from Figure 2 by the total population sizes gives (0.082 packs/day $\times$ 254 million adults + 0.006 packs/day $\times$ 25 million youth) $\times$ 365 days/year $\approx$ 7.7 billion packs. This is 64 percent of the 12 billion packs sold in 2018 as reported in Figure 1. This 64 percent ratio is consistent with the public health literature on under-reported smoking prevalence in sample surveys: for example, Liber and Warner (2018) find 61 percent ratio in the NHIS and about 70 percent in the NSDUH.

For e-cigarettes, multiplying 2018 average vaping for adults and youths from Figure 2 by total population sizes gives (0.03 $\times$ 254 million adults + 0.06 $\times$ 25 million youth) $\times$ 0.58 ml/day $\times$ \$3.90/ml $\approx$ \$7.54 billion. This is nine percent larger than the \$6.9 billion in vapor products sold in 2018 as reported in Figure 1.

B.3 Other Data

E-cigarette User Survey:

- Weight construction. We construct weights using Entropy Weight Rebalancing (Hainmueller 2012), targeting the distribution of gender, income, and e-cigarette use from adults in the sample of BRFSS and the NHIS who report non-zero vaping.
- E-liquid use per day. Several participants record more than 3 ml per day of e-liquid use. We drop their reports from the data, since these are unrealistically large, and winsorize other reports at 2 ml per day.
- Price per day. We construct the weighted mean among participants who report using 3 ml or less e-liquid per day.

E-cigarette Expert Survey:

- Internalities. One expert reports an "infinite" internality of e-cigarettes. We recode this observation as the maximum among experts who report less than an infinite internality.

E-cigarette Tax Rates:

- We use January 1, 2018 tax rates from Tax Foundation (2018). We convert specific taxes to ad valorem taxes using the mean e-cigarette size from RMS and price from the E-cigarette User Survey.

C Price Elasticity Appendix

Table A2: **Own- and Cross-Price Elasticity of Demand for Cigarettes (UPC-level estimates)**

(a) **First Stage and Reduced Form**

	(1)	(2)	(3)
Dependent variable:	ln(cig price)	ln(e-cig price)	ln(cig units)
ln(cig % tax rate + 1)	1.073	-0.130	-1.037
	(0.024)	(0.150)	(0.242)
ln(e-cig % tax rate + 1)	-0.002	0.570	-0.030
	(0.018)	(0.109)	(0.173)
Observations	1,949,823	1,949,823	1,949,823

(b) **Instrumental Variables Estimates**

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	ln(cig units)	ln(cig units)	ln(cig units)	ln(cig units)	ln(cig units)	ln(cig units)
ln(cig price)	-0.974	-6.060	-0.558	-1.321	-1.333	-0.993
	(0.194)	(2.762)	(0.758)	(0.280)	(0.282)	(0.224)
ln(e-cig price)	-0.056	2.090	0.678	0.843	0.771	-0.193
	(0.296)	(0.985)	(0.437)	(0.322)	(0.260)	(0.319)
UPC-cluster FE	Yes	No	Yes	Yes	Yes	Yes
UPC-month FE	Yes	No	No	Yes	Yes	Yes
Division-month FE	Yes	No	No	No	Yes	Yes
Cluster × month trend	Yes	No	No	No	No	Yes
Time-varying state controls	Yes	Yes	Yes	Yes	Yes	No
Observations	1,949,823	1,952,925	1,949,875	1,949,823	1,949,823	1,949,823

Notes: This table presents estimates of the own- and cross-price elasticity of demand for cigarettes from an analogue to Equation (16), using UPC-cluster-month data. There are 51 geographic clusters: the two counties that have e-cigarette taxes, each of the contiguous 48 states (excluding those two counties), and Washington, D.C. Standard errors are clustered by geographic cluster. Observations are weighted by the UPC's sales in other clusters in that calendar year, divided by total sales in other clusters across all UPCs in that year. Panel (a) presents the first stage and reduced form, using the same set of controls as in our primary estimate in column 1 of Panel (b). Panel (b) presents the instrumental variables estimates. Time-varying state controls are the state unemployment rate and beer tax rate as well as indicators for whether the state has an indoor vaping ban, has a medical marijuana law, has passed or implemented a prescription drug program, and implemented the Medicaid expansion.

Table A3: **Own- and Cross-Price Elasticity of Demand for E-cigarettes, Robustness**

Dependent variable: ln(e-cig units)	(1) 18-month window	(2) Exclude 1(quarter of e-cig tax) controls	(3) Exclude imputed volumes	(4) Exclude specific-tax clusters
ln(e-cig price)	-1.137	-1.154	-1.297	-1.276
	(0.455)	(0.544)	(0.505)	(0.514)
ln(cig price)	0.405	0.442	0.443	0.444
	(0.574)	(0.593)	(0.610)	(0.562)
Observations	499,664	499,664	496,070	457,997

Notes: This table presents instrumental variables estimates of the own- and cross-price elasticity of demand for e-cigarettes from Equation (16), using UPC-cluster-month data. There are 51 geographic clusters: the two counties that have e-cigarette taxes, each of the contiguous 48 states (excluding those two counties), and Washington, D.C. Standard errors are clustered by geographic cluster. Observations are weighted by the UPC's sales in non-taxed clusters in that calendar year, divided by total sales across all UPCs in that year in non-taxed clusters. All columns include UPC-cluster, UPC-month, and Census division-month fixed effects and cluster-specific linear time trends, as well as time-varying state controls. Column 1 includes additional controls so as to identify the elasticities only using an 18-month window around e-cigarette tax changes. Column 2 excludes controls for the quarter of the e-cigarette tax treatment and interaction with the e-cigarette tax, to identify the elasticities also off of the effects in the first three months after a tax change. Column 3 excludes e-cigarette UPCs with imputed volumes. Column 4 excludes states that ever have a specific (i.e. per-milliliter) e-cigarette tax.

Table A4: **Own- and Cross-Price Elasticity of Demand for Cigarettes, Robustness**

Dep. variable: ln(cig units)	(1) 18-month window	(2) Exclude 1(quarter of e-cig tax) controls	(3) Exclude specific-tax states
ln(cig price)	-0.974	-0.969	-0.978
	(0.194)	(0.186)	(0.187)
ln(e-cig price)	-0.066	-0.072	-0.113
	(0.279)	(0.312)	(0.192)
Observations	1,949,823	1,949,823	1,764,557

Notes: This table presents instrumental variables estimates of the own- and cross-price elasticity of demand for cigarettes from an analogue to Equation (16), using UPC-cluster-month data. There are 51 geographic clusters: the two counties that have e-cigarette taxes, each of the contiguous 48 states (excluding those two counties), and Washington, D.C. Standard errors are clustered by geographic cluster. Observations are weighted by the UPC's sales in other clusters in that calendar year, divided by total sales in other clusters across all UPCs in that year. All columns include UPC-cluster, UPC-month, and Census division-month fixed effects and cluster-specific linear time trends, as well as time-varying state policy controls. Column 1 includes additional controls so as to identify the elasticities only using an 18-month window around e-cigarette tax changes. Column 2 excludes controls for the quarter of the e-cigarette tax treatment and interaction with the e-cigarette tax, to identify the elasticities also off of the effects in the first three months after a tax change. Column 3 excludes states that ever have a specific (i.e. per-milliliter) e-cigarette tax.

Figure A2: **Event Study of E-cigarette Tax Changes without Linear Time Trends**

(a) **First Stage: Effect on ln(Price)**

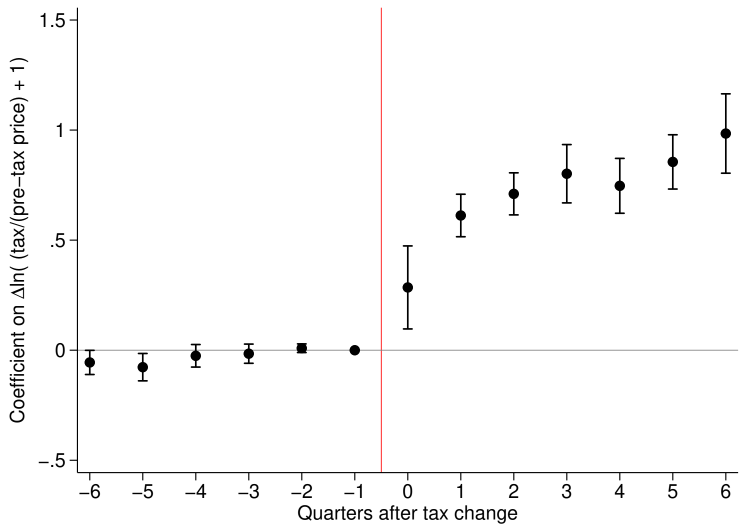

(b) **Reduced Form: Effect on ln(Quantity Sold + 1)**

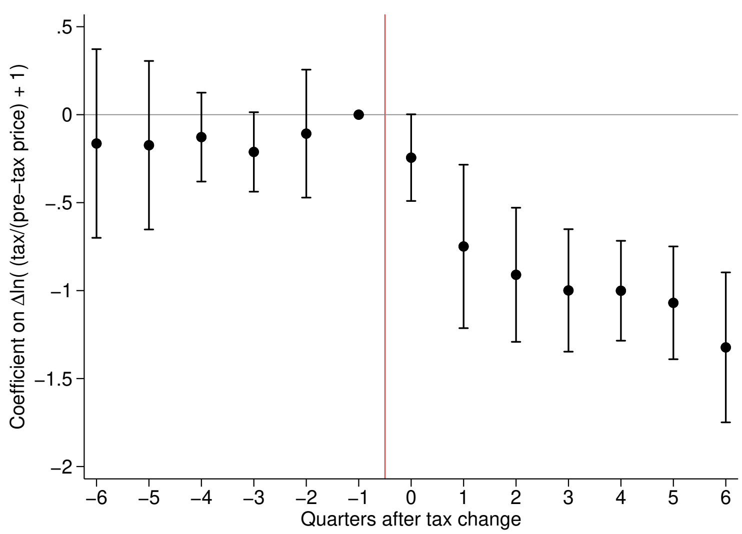

Notes: This figure presents estimates of the η_q parameters from Equation (17), an event study of the effects of e-cigarette tax changes. This parallels Figure 3, except it excludes cluster-specific linear time trends. Panel (a) presents the first stage regression of ln(e-cigarette price) on the change in the log tax variable. Panel (b) presents the reduced form regression of the ln(e-cigarette units sold) on the change in the log tax variable.

Figure A3: **Event Study of E-cigarette Tax Changes on Cigarette Demand**

(a) **With Cluster-Specific Linear Time Trends**

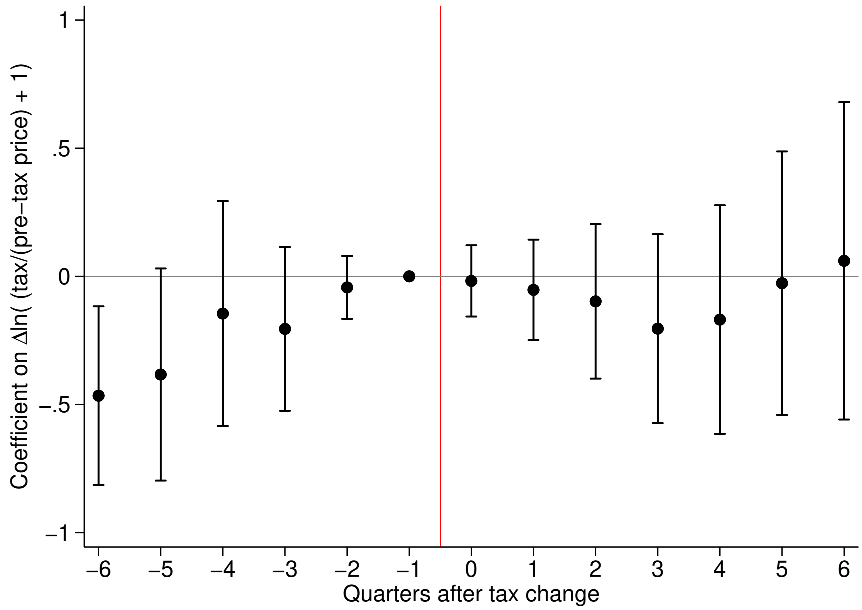

(b) **Without Cluster-Specific Linear Time Trends**

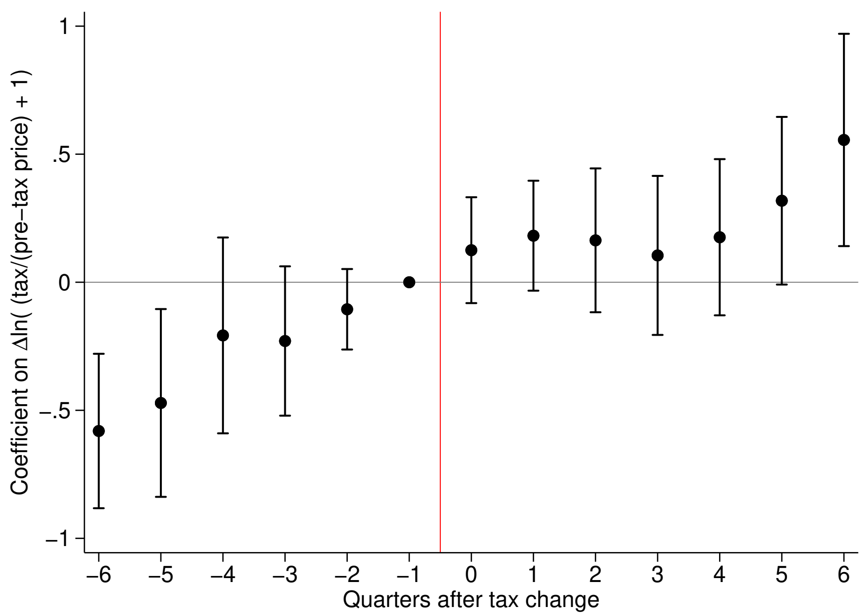

Notes: This figure presents estimates of the η_q parameters from Equation (17), an event study of the effects of e-cigarette tax changes. This parallels Figure 3, except with combustible cigarette purchases as the dependent variable. Panel (a) presents estimates with cluster-specific linear time trends. Panel (b) presents estimates without cluster-specific linear time trends.

D Substitution Patterns Appendix

Figure A4: **Demographic Predictors of Vaping, by Dataset**

(a) **Adults**

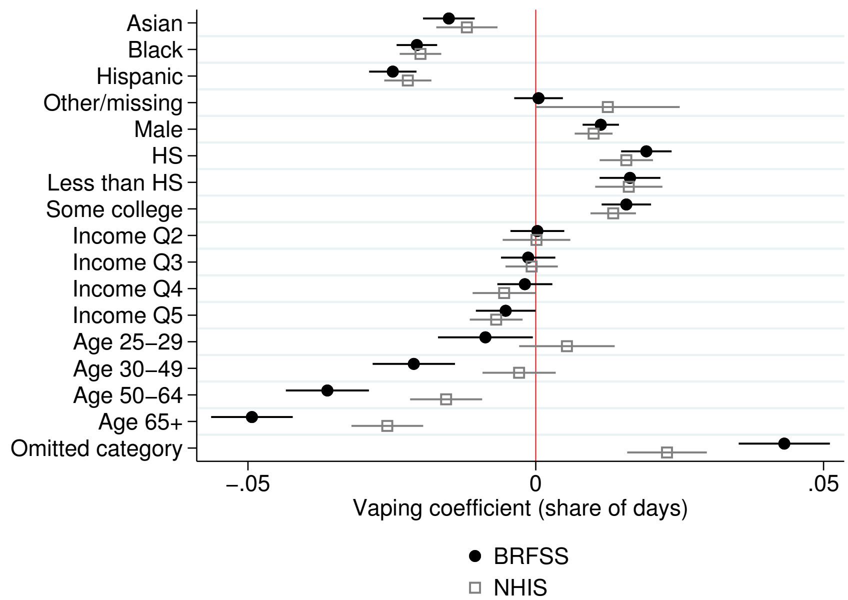

(b) **Youth**

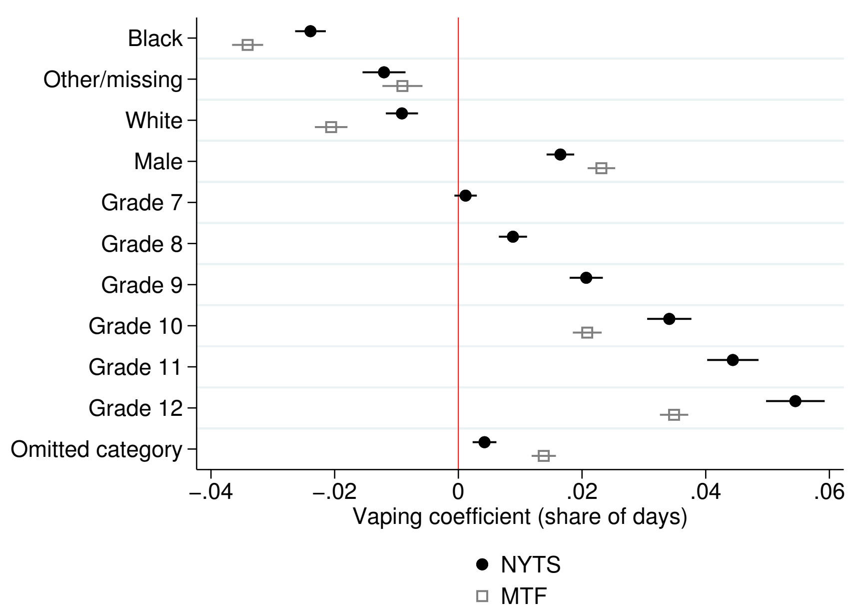

Notes: These figures present coefficients from Equation (20), a regression of vaping on demographic indicators, estimated separately by dataset. For adults, the omitted categories are White, female, college graduate, the lowest income quintile, and age group 18-24. For youth, the omitted categories are White, female, and grade 6. Panel (a) pools 2016–2018 data from BRFSS and NHIS; Panel (b) pools 2014–2018 data from MTF and NYTS. Standard errors are clustered by demographic cell.

Figure A5: **Demographic Predictors of E-cigarette, Social Media, and Internet Use**

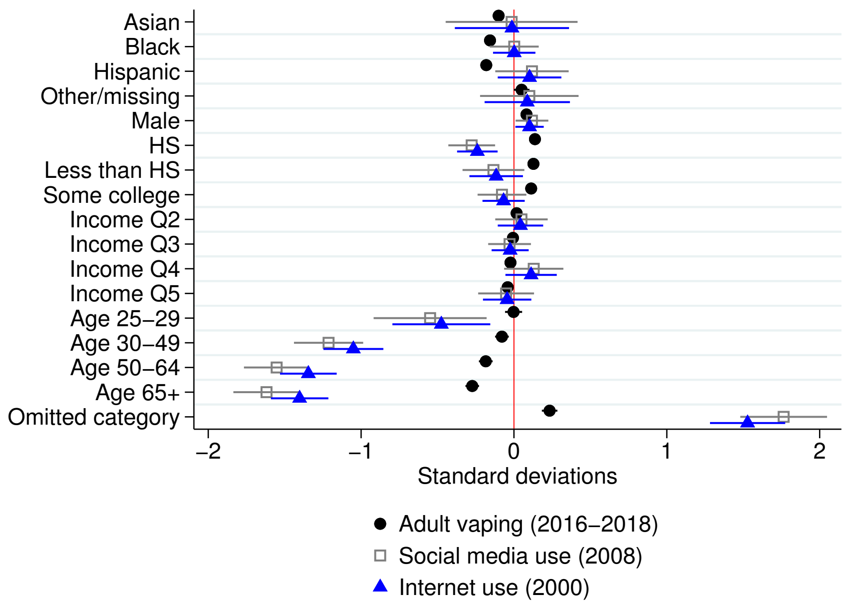

Notes: These figures present coefficients from regressions of vaping, social media use, or internet use on demographic indicators. Each dependent variable is normalized into standard deviation units for comparability. For adults, the omitted categories are White, female, college graduate, the lowest income quintile, and age group 18-24. For youth, the omitted categories are White, female, and grade 6. Standard errors are clustered by demographic cell.

Figure A6: **Smoking and Vaping Trends by Sex**

(a) **Adults**

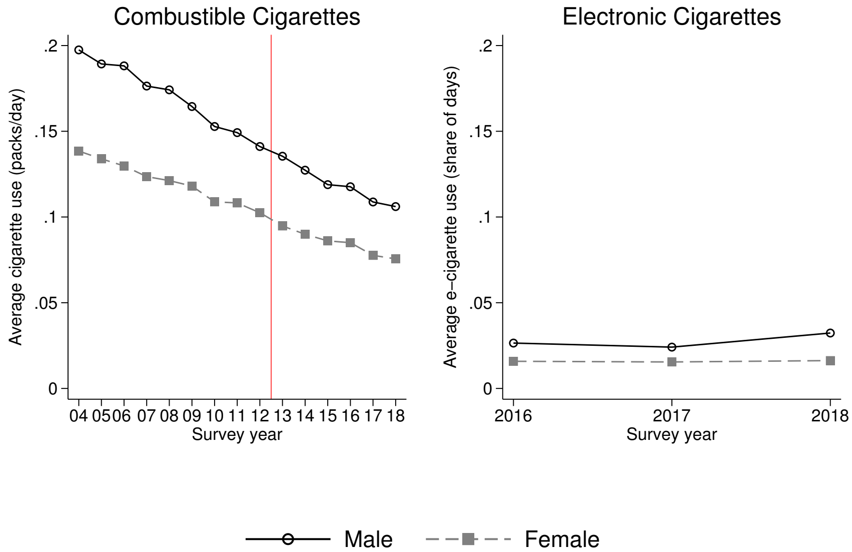

(b) **Youth**

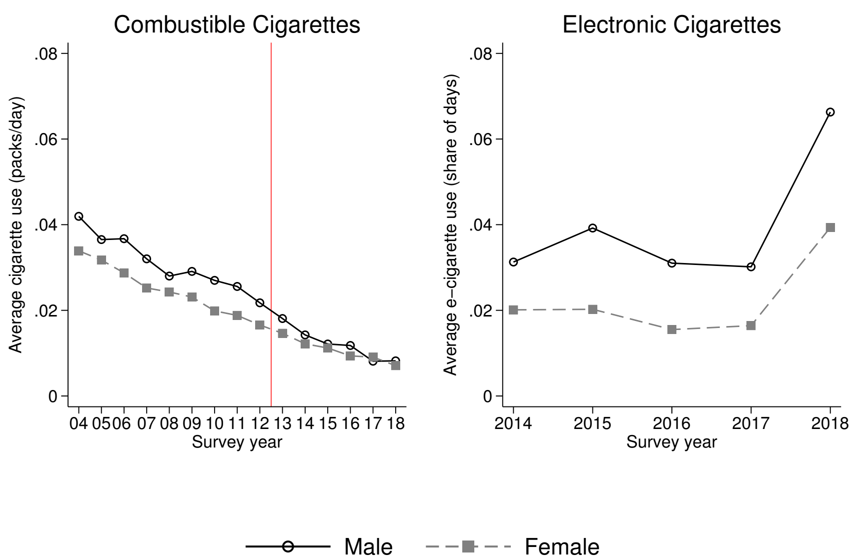

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

Figure A7: **Smoking and Vaping Trends by Race/Ethnicity**

(a) **Adults**

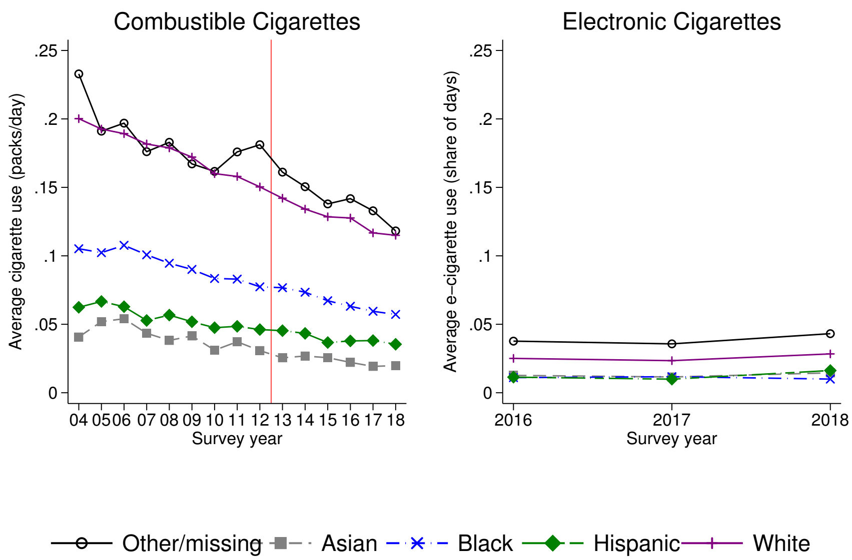

(b) **Youth**

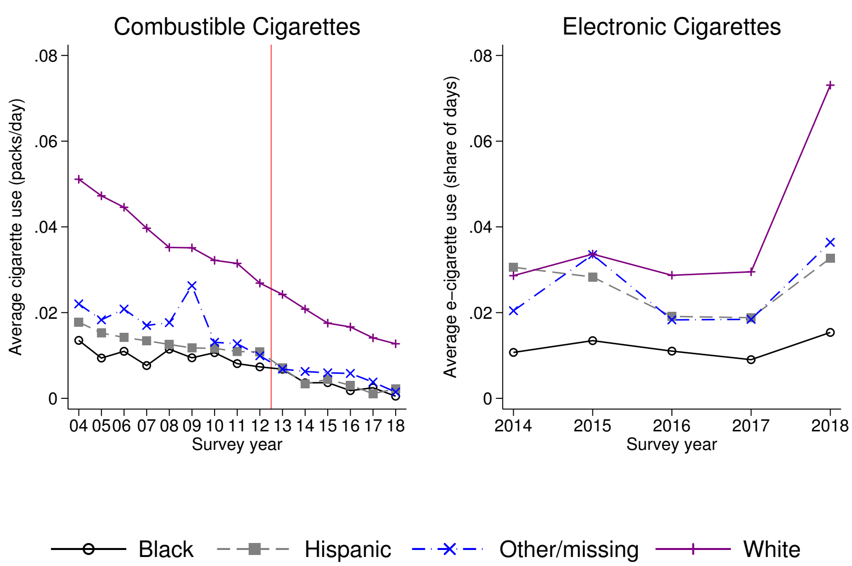

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

Figure A8: **Smoking and Vaping Trends by Age/Grade**

(a) **Adults**

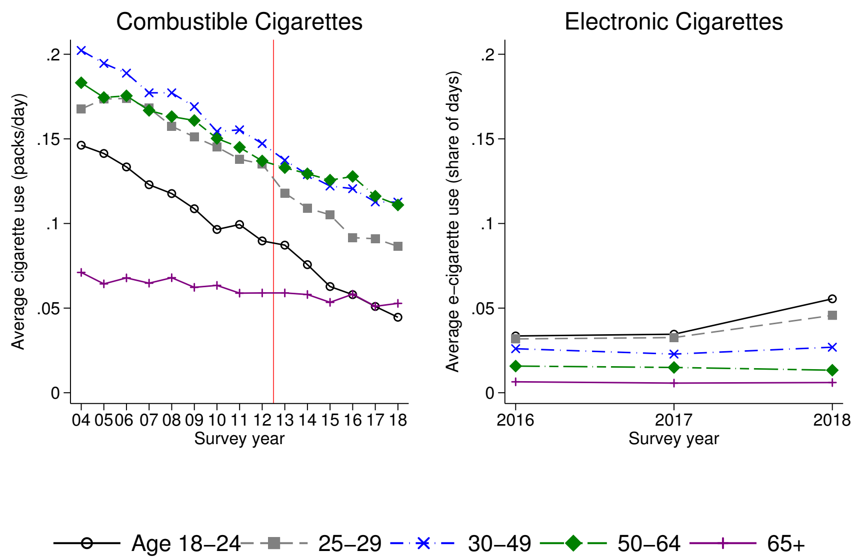

(b) **Youth**

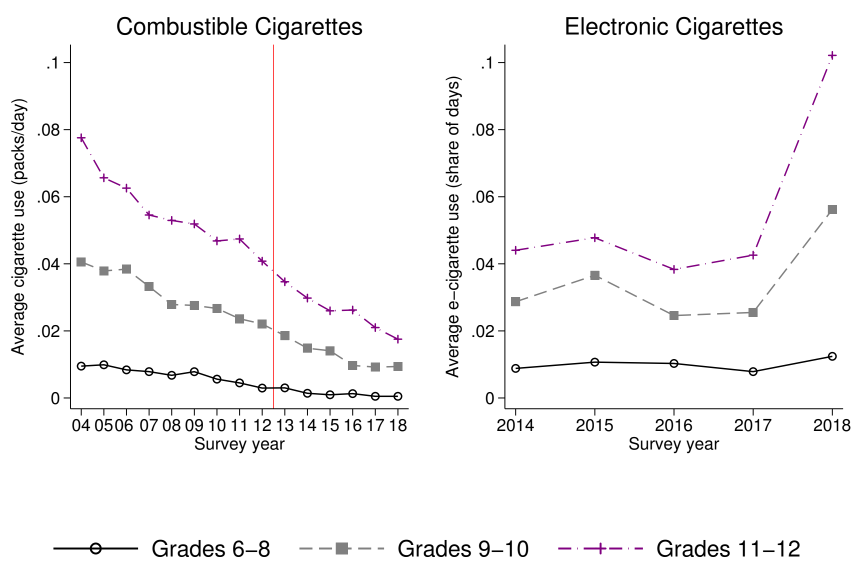

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

Figure A9: **Smoking and Vaping Trends by Education and Income, for Adults**

(a) **Education**

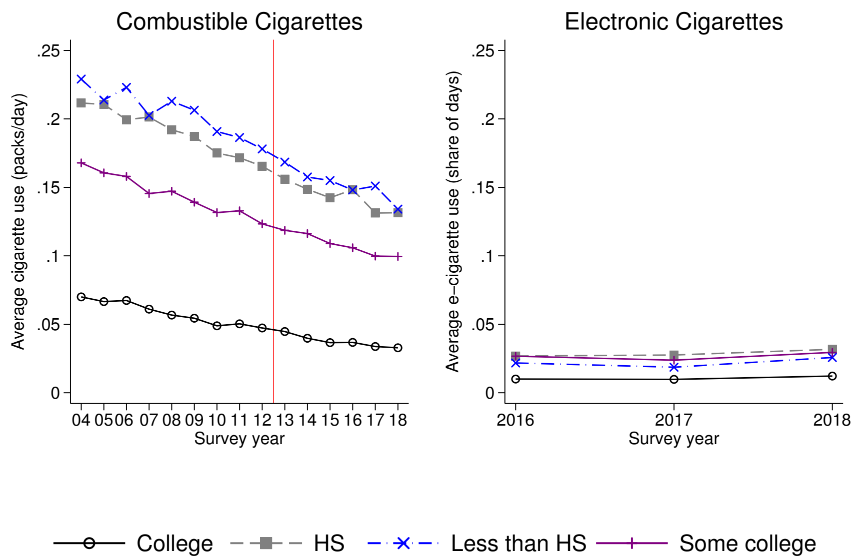

(b) **Income**

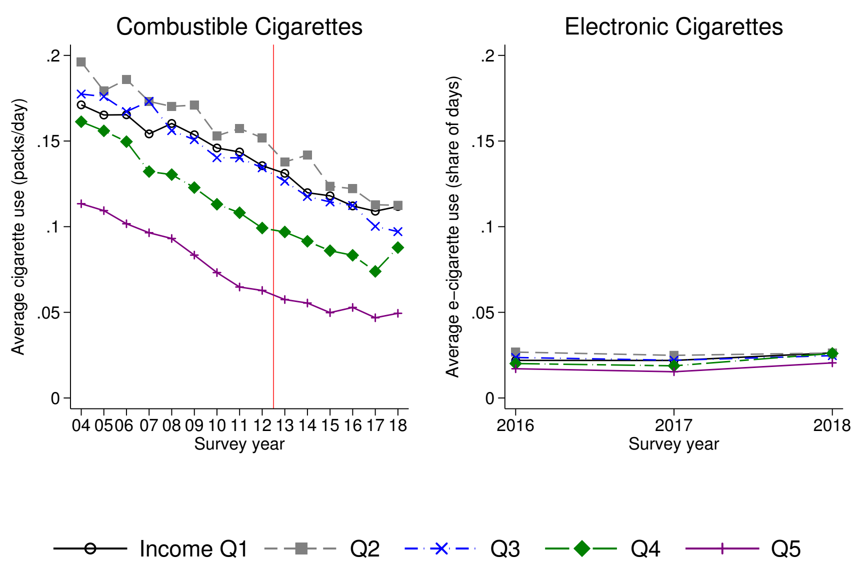

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

Figure A10: **Difference in Smoking Trends by Sex**

(a) **Adults**

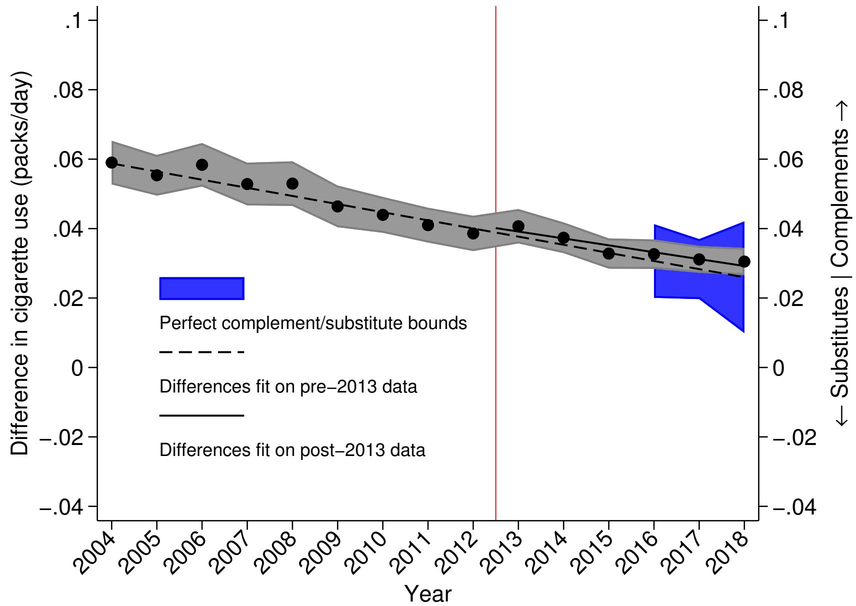

(b) **Youth**

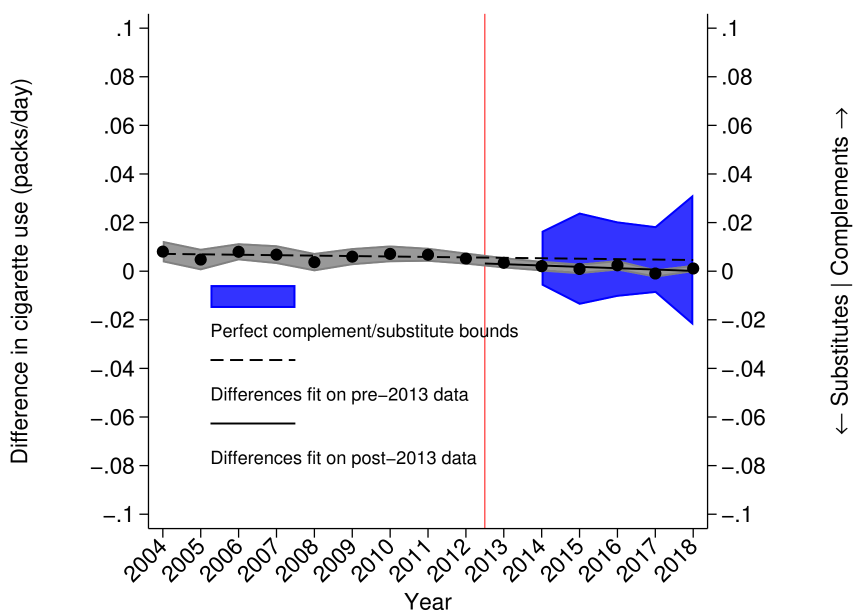

Notes: These figures present the difference in cigarette use for men versus women. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. The perfect complement (substitute) bounds show the difference in cigarette use that would be expected if daily vaping were a perfect complement (substitute) for smoking one pack of cigarettes per day. To construct perfect complement (substitute) bounds, we predict the difference in cigarette use using the pre-2013 time trend, then add (subtract) the difference in share of days vaped.

Figure A11: **Difference in Smoking Trends by Race**

(a) **Adults**

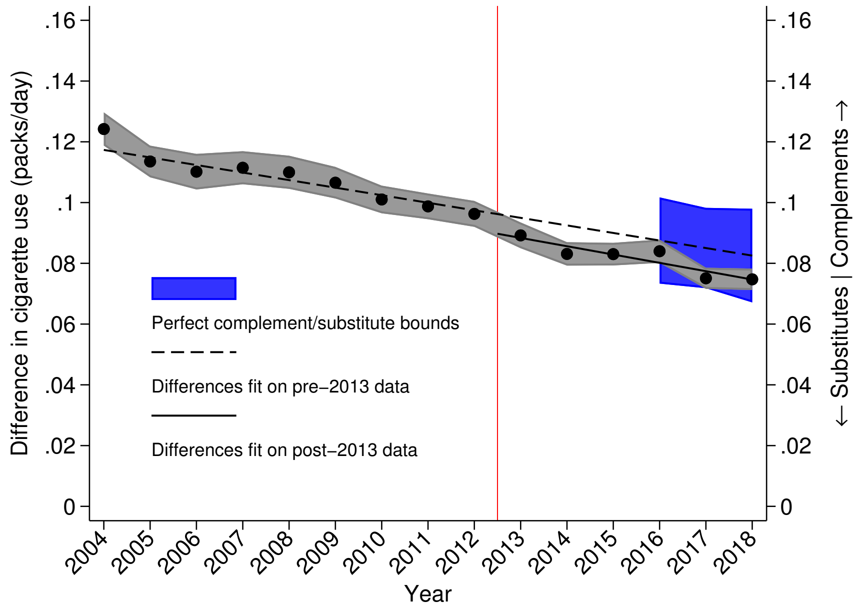

(b) **Youth**

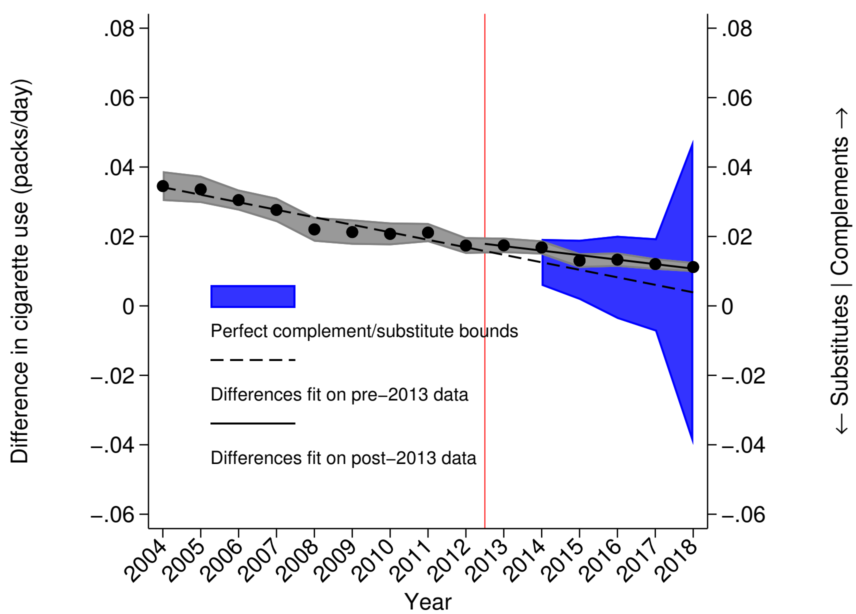

Notes: These figures present the difference in cigarette use for Whites and other races versus non-Whites (for adults) and Whites versus non-Whites (for youth). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. The perfect complement (substitute) bounds show the difference in cigarette use that would be expected if daily vaping were a perfect complement (substitute) for smoking one pack of cigarettes per day. To construct perfect complement (substitute) bounds, we predict the difference in cigarette use using the pre-2013 time trend, then add (subtract) the difference in share of days vaped.

Figure A12: **Difference in Smoking Trends by Age/Grade**

(a) **Adults**

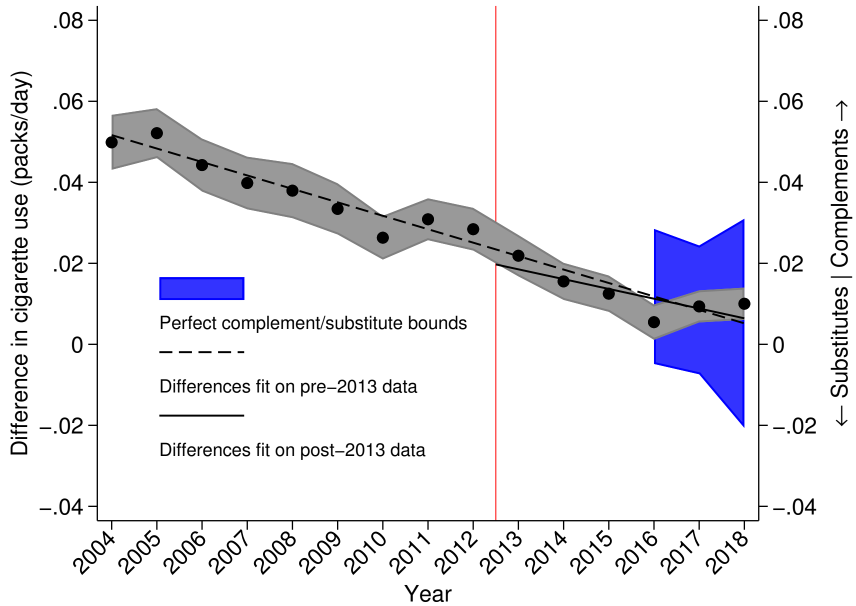

(b) **Youth**

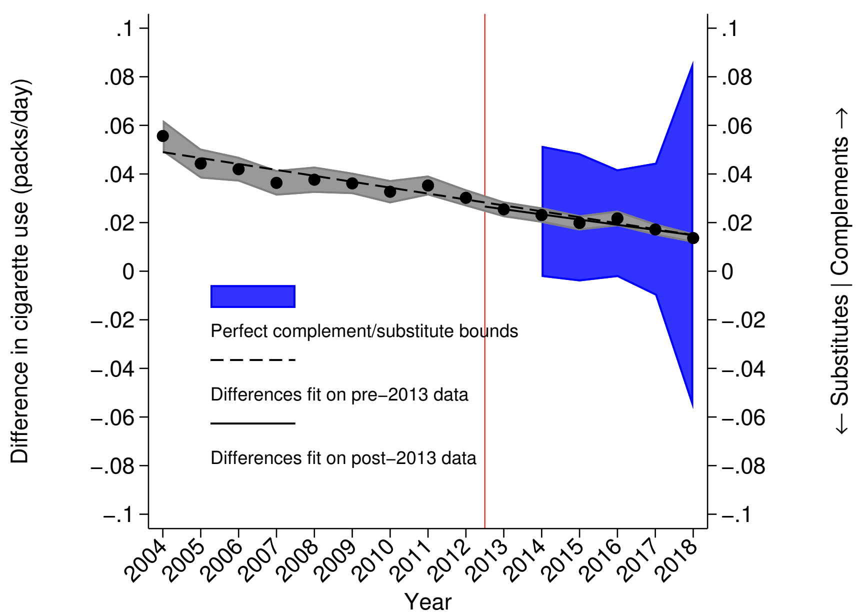

Notes: These figures present the difference in cigarette use by year for age ≤ 49 versus age ≥ 50 (for adults) and for grades ≥ 11 versus grades ≤ 10 (for youth). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. The perfect complement (substitute) bounds show the difference in cigarette use that would be expected if daily vaping were a perfect complement (substitute) for smoking one pack of cigarettes per day. To construct perfect complement (substitute) bounds, we predict the difference in cigarette use using the pre-2013 time trend, then add (subtract) the difference in share of days vaped.

Figure A13: **Difference in Smoking Trends by Education and Income, for Adults**

(a) **Education**

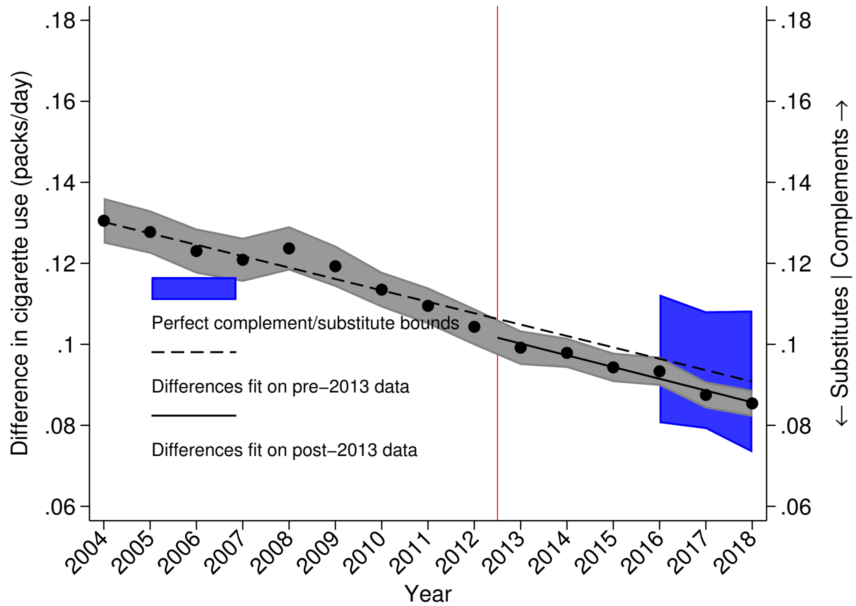

(b) **Income**

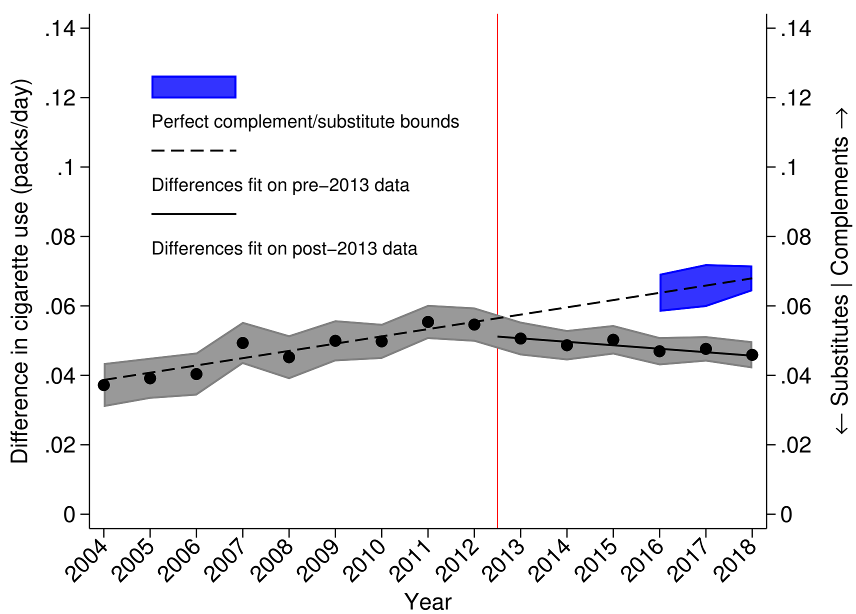

Notes: These figures present the difference in cigarette use by year for adults without versus with college degrees (Panel (a)) and adults in the bottom three versus top two income quintiles (Panel (b)). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. The perfect complement (substitute) bounds show the difference in cigarette use that would be expected if daily vaping were a perfect complement (substitute) for smoking one pack of cigarettes per day. To construct perfect complement (substitute) bounds, we predict the difference in cigarette use using the pre-2013 time trend, then add (subtract) the difference in share of days vaped.

D.1 Combined Substitution Estimates

In this appendix, we describe how we form combined estimates of the substitution parameter σ using both the RMS estimates from Section 4 and the sample surveys from Section 5. σ is in units of packs of cigarettes per day vaped. Define

$$\sigma_1 := \frac{\chi^e}{\eta}\frac{\tilde{p}^e\Gamma}{\tilde{p}^c}, \tag{36}$$

where Γ (ml/average day vaped) converts $\tilde{p}^e$ to units of dollars per day vaped. Further define

$$\sigma_{\theta 2} := \frac{\chi^c}{\eta}\frac{q^c_\theta}{q^e_\theta} \tag{37}$$

and note that $\hat{\sigma}_{\theta 2}$ is already in units of packs per day. The empirical estimates are the respective plug-in estimators using $\hat{\chi}^e$, $\hat{\chi}^c$, and $\hat{\eta}$ from Table 2 and A2, and $\hat{q}^j_\theta$, $\hat{\tilde{p}}^j$ and $\hat{\Gamma}$ from Table 4 for $j \in \{c, e\}$. We form one estimate of $\hat{\sigma}_1$ using the primary estimate from Table (2) (Panel B, Column 1), and a second estimate of $\hat{\sigma}_1$ using the estimates of $\hat{\chi}^e$ and $\hat{\eta}$ estimated without cluster-specific linear trends (Column 5). We form standard errors on $\hat{\sigma}_1$ and $\hat{\sigma}_2$ using the delta method; the variance-covariance matrix is diagonal except for the covariance term between $\hat{\eta}$ and χ^e.

We combine $\hat{\sigma}_1$ and $\hat{\sigma}_2$ using Classical Minimum Distance (CMD) using:

$$\begin{bmatrix} 1 \\ 1 \end{bmatrix} \sigma_\theta - \begin{bmatrix} \sigma_1 \\ \sigma_{\theta 2} \end{bmatrix} = \mathbf{0} \ , \tag{38}$$

noting that

$$\begin{bmatrix} \sigma_1 \\ \sigma_{\theta 2} \end{bmatrix} \sim N\left(0, \begin{bmatrix} s_1^2 & s_{12} \\ s_{12} & s_1^2 \end{bmatrix}\right). \tag{39}$$

We use $\hat{s}_1^2$ and $\hat{s}_2^2$ from the initial delta method estimation. We estimate s_{12} as follows:

$$\begin{aligned} s_{12} &:= Cov\left(\frac{\chi^c}{\eta}\frac{q^c_\theta}{q^e_\theta}, \frac{\chi^e}{\eta}\frac{\tilde{p}^e}{\tilde{p}^c}\Gamma\right) && (40)\\ &= \chi^c \frac{q^c_\theta}{q^e_\theta}\frac{\tilde{p}^e}{\tilde{p}^c}\Gamma Cov\left(\frac{1}{\eta}, \frac{\chi^e}{\eta}\right) && (41)\\ &\approx \chi^e\chi^c \frac{q^c_\theta}{q^e_\theta}\frac{\tilde{p}^e}{\tilde{p}^c}\Gamma V\left(\frac{1}{\eta}\right) && (42) \end{aligned}$$

where the second line follows since the parameters taken outside the covariance are all estimated from separate datasets, and we assume that the covariance between χ^e and $1/\eta$ is small. We estimate $V\left(\frac{1}{\eta}\right)$ from the delta method, and form $\hat{s}_{12}$ using a plug-in estimator.

We also combine $\hat{\sigma}_1$ and $\hat{\sigma}_2$ with our estimates from Section 5 using CMD. Table A5 presents our results.

Table A5: **Estimates of Substitution Parameter σ**

	(1) E-cig cross-price elasticity	(2) E-cig cross-price elasticity (no trends)	(3) Cig cross-price elasticity	(4) Combined RMS	(5) Demo. analysis	(6) Combined RMS and demo.
Adult σ	-0.059	-0.191	0.144	-0.056	0.035	-0.020
	(0.091)	(0.113)	(0.766)	(0.090)	(0.112)	(0.070)
Youth σ	-0.059	-0.191	0.005	-0.000	0.013	0.007
	(0.091)	(0.113)	(0.027)	(0.026)	(0.022)	(0.017)

Notes: This table presents estimates of the substitution parameter σ for youth and adults. Column 1 presents $\hat{\sigma}$ from Equation (36) using our primary $\hat{\eta}$ and $\hat{\chi}^e$ from Table 2 (Panel (b), column 1). Column 2 presents $\hat{\sigma}$ from Equation (36) using $\hat{\eta}$ and $\hat{\chi}^e$ estimated without cluster-specific linear trends (Table 2, panel (b), column 5). Column 3 presents $\hat{\sigma}$ from Equation (37) using $\hat{\chi}^c$ from Appendix Table A2 (Panel (b), column 1). Column 4 combines the estimates in columns 1 and 3 using Equation (38). Column 5 re-states estimates from the demographic shift-share analysis in Section 5. Column 6 combined estimates from columns 4 and 5 using Classical Minimum Distance.

D.2 Marijuana Use

We study the time series of teen marijuana use during the period that e-cigarettes use became common among teens using the MTF. A concern about our welfare analysis is that we do not account for substitution from e-cigarettes into possibly harmful drugs like marijuana; there is a particular concern that vaping technologies make it easier to vape marijuana. In this section, we provide evidence against this concern by documenting no change in *aggregate* marijuana consumption over this time period; while *vaping* marijuana becomes more popular, *total* marijuana use exhibits a small decline.

Marijuana use in the MTF. We focus on youth vaping, for whom the concerns about substitution into marijuana products are most salient. The MTF provides several measures of marijuana use. First, beginning in 2014, the MTF asks respondents the number of times they consumed marijuana last year in any form. Second, beginning in 2017, the MTF asks respondents the number of times that they consumed marijuana last month in any form. Third, beginning in 2017, the MTF asks respondents the number of times that they vaped marijuana last month. We standardize these variables to construct the number of times the respondent consumed vaped marijuana each day. Due to interval censoring and top coding, the marijuana consumption measures do not align perfectly. In particular, both the monthly and annual marijuana measures are subject to significant top coding; the participant cannot report consuming marijuana more than 40 times in the past month or year. As a result, the annual measure naturally lies below the monthly estimate. However, we are concerned with trends in marijuana use as e-cigarette use becomes popular and simply discuss changes in marijuana use, comparing each measure over time.

Results. In Appendix Figure A14, panel (a), we present the time series of e-cigarette use against the time series of our three measures of marijuana use; panel (b) focuses on grades 11–12, which has higher rates of both e-cigarette use and marijuana consumption. This figure illustates that while *vaping* marijuana does become more popular in 2018 (as e-cigarette use grew), the time series of *aggregate* marijuana use exhibits no change over this period. In fact, the monthly measure of marijuana consumption shows a small decline from 2017–2018 in both the full sample and grades 11–12. While we do not conduct a full substitution analysis, these figures suggest that the aggregate data are inconsistent with the concern that our welfare analysis neglects important distortions induced by e-cigarette use.

Figure A14: **Trends in Youth Marijuana Use**

(a) **All**

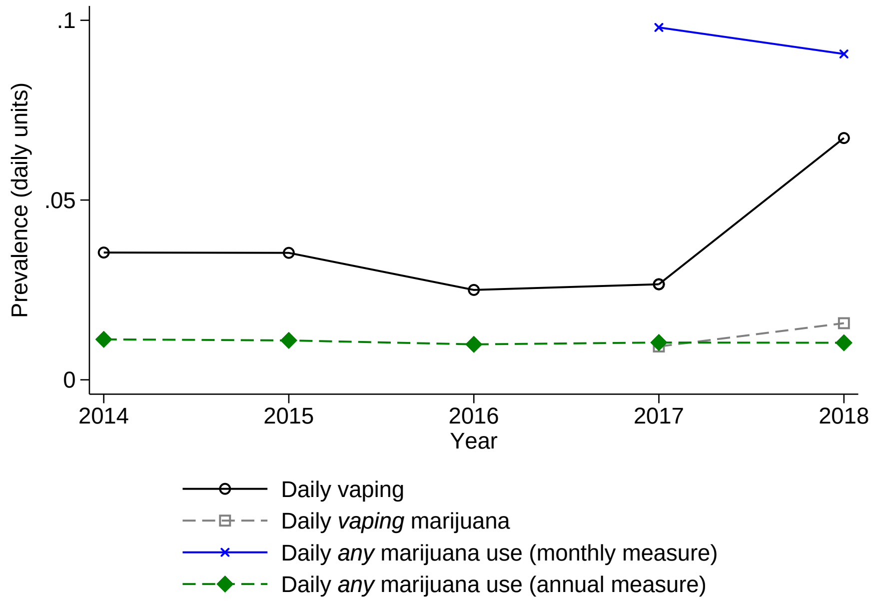

(b) **Grades 11–12**

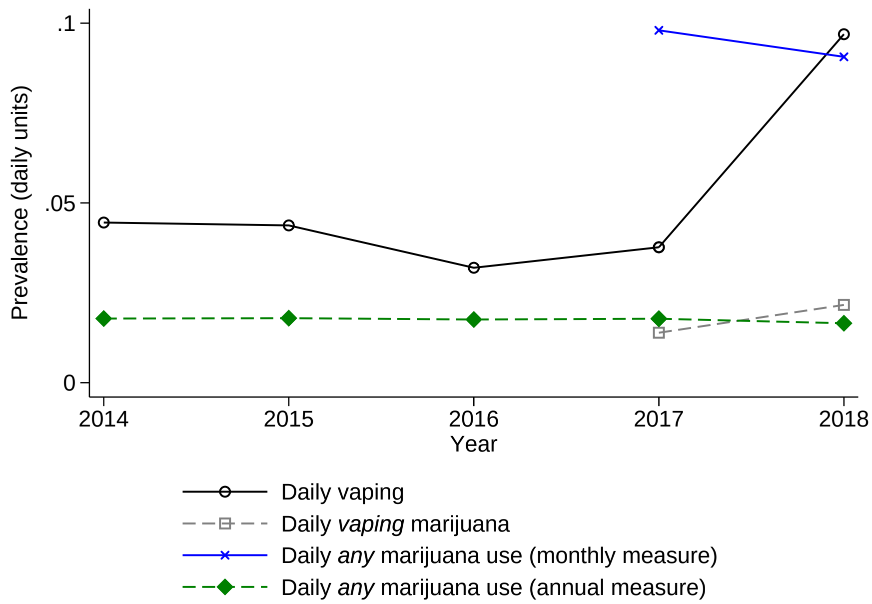

Notes: This figure presents trends in marijuana and e-cigarette use in the Monitoring the Future (MTF) survey. Panel (a) presents the full sample, while panel (b) focuses on grades 11 and 12. The black lines present our daily vaping measure. The gray lines present the average daily *vaping* marijuana use, constructed from an MTF question that asks about the number of times the respondent vaped in the past month. The blue line presents the average daily marijuana consumption of any form, constructed from an MTF question that asks about the number of times the respondent consumed marijuana in the past *month*. The green line presents the same measure, but from an MTF question that asks about the number of times the respondent consumed marijuana in the past *year*. The green line lies below the blue line due to top-coding.

E Expert Survey Appendix

Figure A15: **Expert Survey: Responses from Public Health Researchers and Economists**

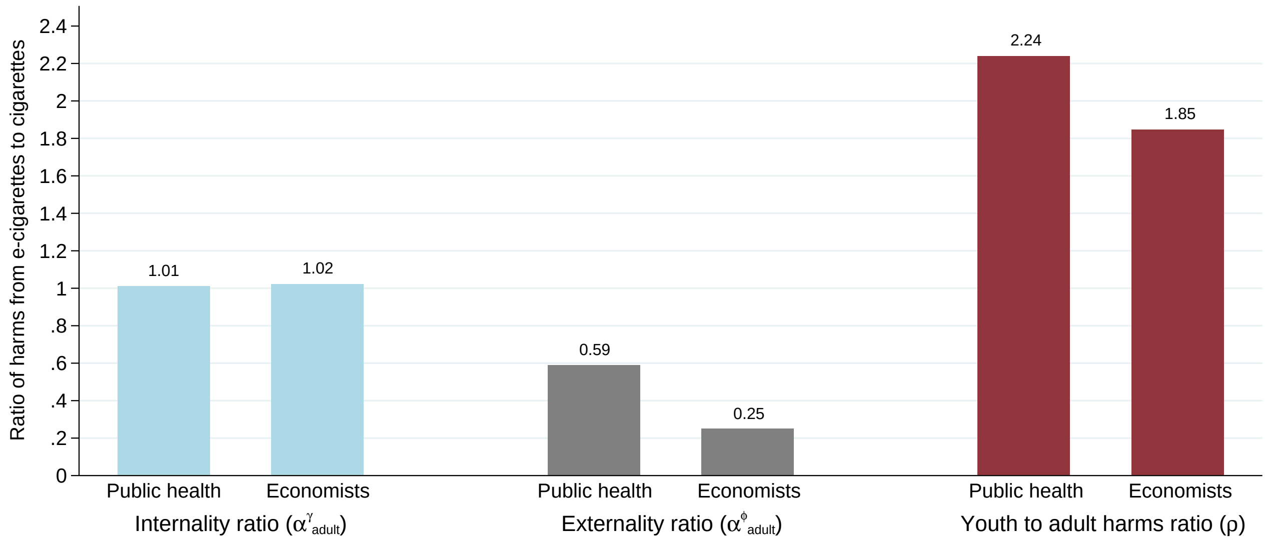

Notes: In our expert survey, we elicited the ratios of internalities and externalities from vaping relative to smoking and the ratio of uninternalized harms from vaping for youth relative to adults. This figure presents the averages of those ratios separately for public health researchers cited in National Academy of Science (2018) and economists.

Figure A16: **Expert Survey: Uninternalized Harms from Vaping for Youth Relative to Adults**

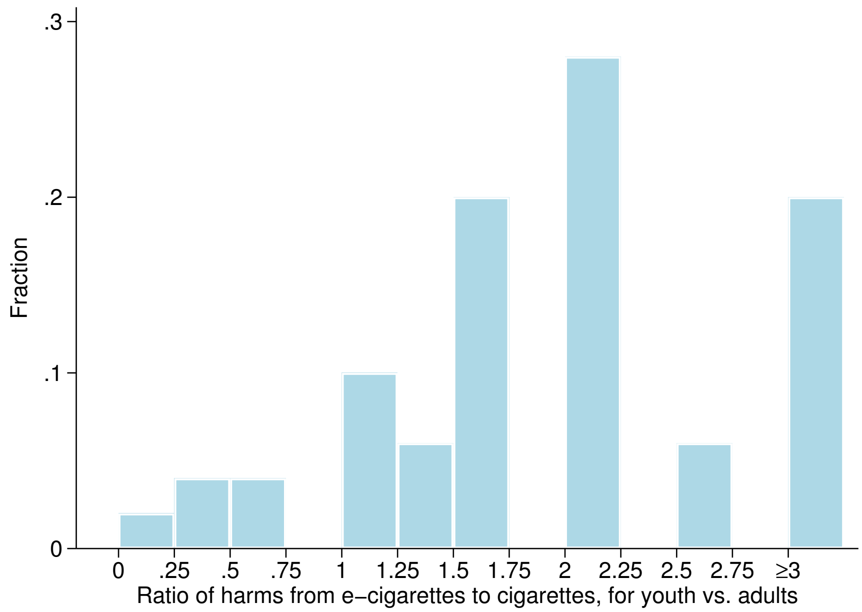

Notes: In our expert survey, we elicited the ratio of uninternalized harms from vaping for youth relative to adults. This figure presents the distributions of that ratio across experts.

F Welfare Analysis Appendix

The version of Equation (14) for empirical implementation is

$$\tau^{e*} = \frac{\sum_{\Theta} s_\theta q_\theta^e \Gamma \left[\varphi_\theta^e + (\sigma_\theta/\Gamma)\left(\varphi_\theta^c - \tau^c\right)\right]}{\sum_{\Theta} s_\theta q_\theta^e \Gamma}, \tag{43}$$

where $\varphi_a^e = \alpha^\gamma \frac{\Omega_a}{\Gamma} \cdot \gamma^c + \alpha^\phi \frac{\Omega_a}{\Gamma} \cdot \phi^c$ and $\varphi_y^e = \rho \varphi_a^e$. Vaping quantity q_θ^e is in units of share of days, σ_θ is in units of packs of cigarettes per day vaped, and Γ is in units of ml fluid/day vaped. τ^{e*} and φ_θ^e are in units of \$/ml.

The version of Equation (15) for empirical implementation is

$$\Delta\overline{W} = 365 \times \sum_{\theta\in\{a,y\}} s_\theta \left[\underbrace{q^e_\theta \Gamma \frac{\tilde{p}^e}{-2\eta}}_{\text{perceived CS change}} - \underbrace{\left(-q^e_\theta\Gamma\right)\left(\varphi^e_\theta - \tau^e\right)}_{\text{e-cigarette distortion change}} - \underbrace{q^e_\theta\Gamma\cdot\left(-\sigma_\theta/\Gamma\right)\left(\varphi^c - \tau^c\right)}_{\text{cigarette distortion change}} \right], \tag{44}$$

where $\Delta\overline{W}$ is in units of dollars per person-year.

Figure A17: **Parameter Regions where Youth E-cigarette Ban Increases Welfare**

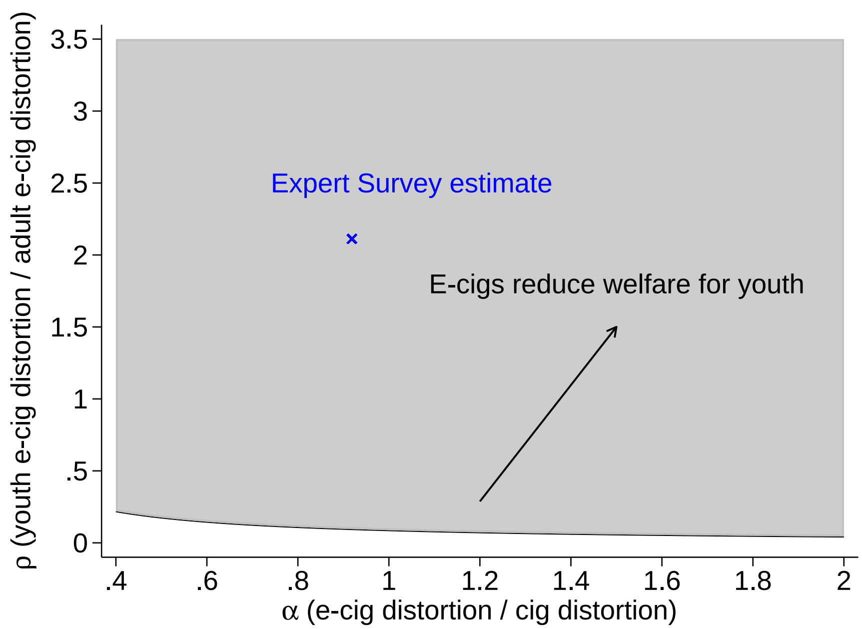

Notes: This figure presents parameter regions where a youth e-cigarette ban increases welfare, using Equation (15). All parameters other than α_{adult} and ρ are set at their means presented in Table 4.