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

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

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**

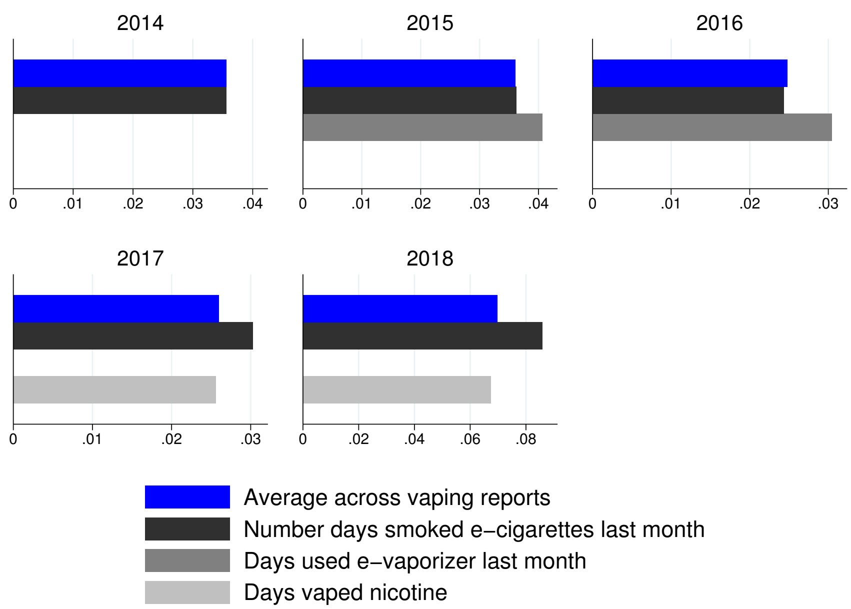

Notes: This figure presents vaping rates by year and question from the Monitoring the Future survey.

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)**

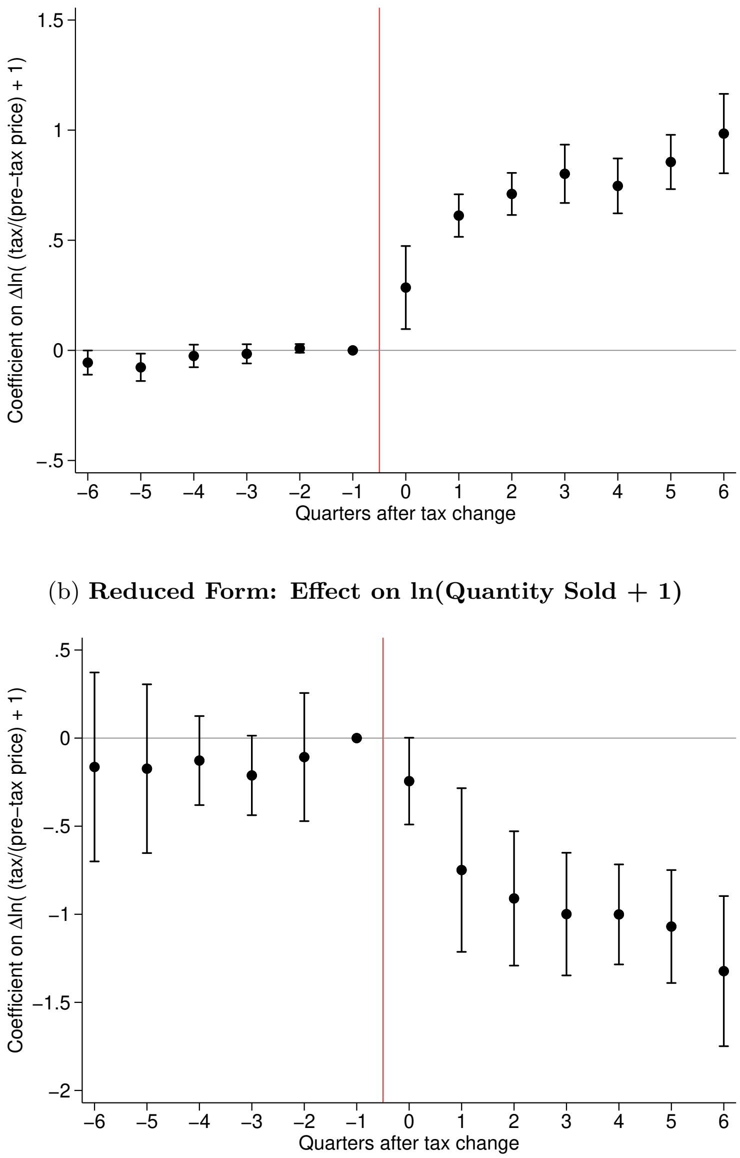

Notes: This figure presents estimates of the η_q parameters from Equation (17), an event study of the effects of e-cigarette tax changes. This parallels Figure 3, except it excludes cluster-specific linear time trends. Panel (a) presents the first stage regression of ln(e-cigarette price) on the change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

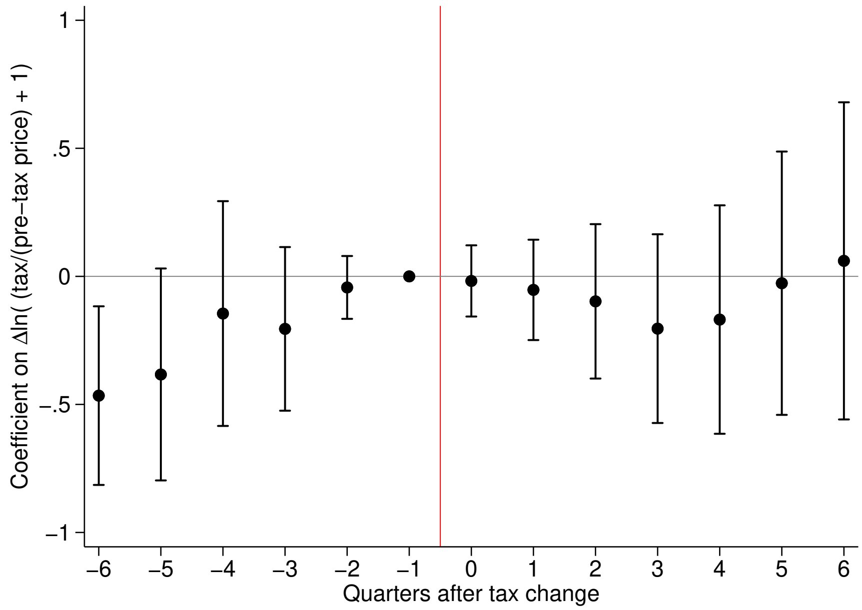

(b) **Without Cluster-Specific Linear Time Trends**

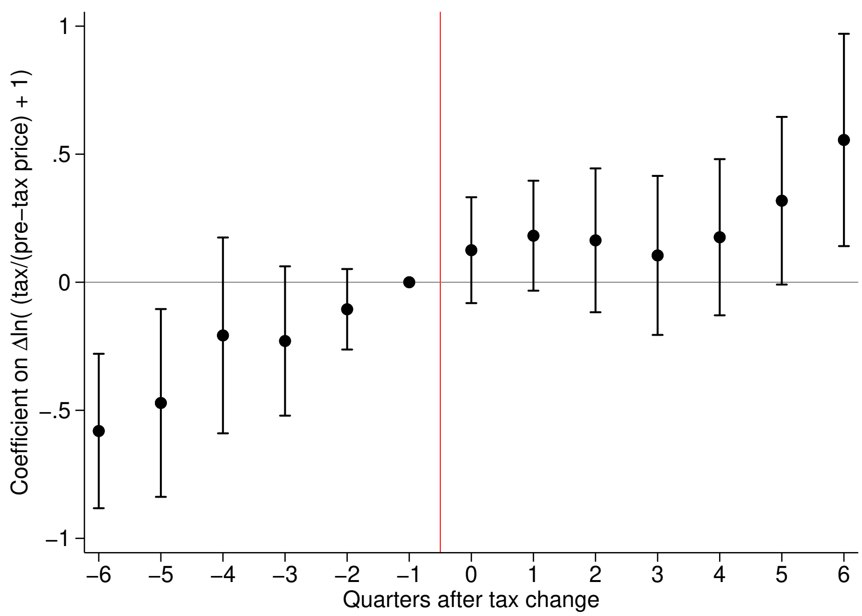

Notes: This figure presents estimates of the η_q parameters from Equation (17), an event study of the effects of e-cigarette tax changes. This parallels Figure 3, except with combustible cigarette purchases as the dependent variable. Panel (a) presents estimates with cluster-specific linear time trends.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

D Substitution Patterns Appendix

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

(a) **Adults**

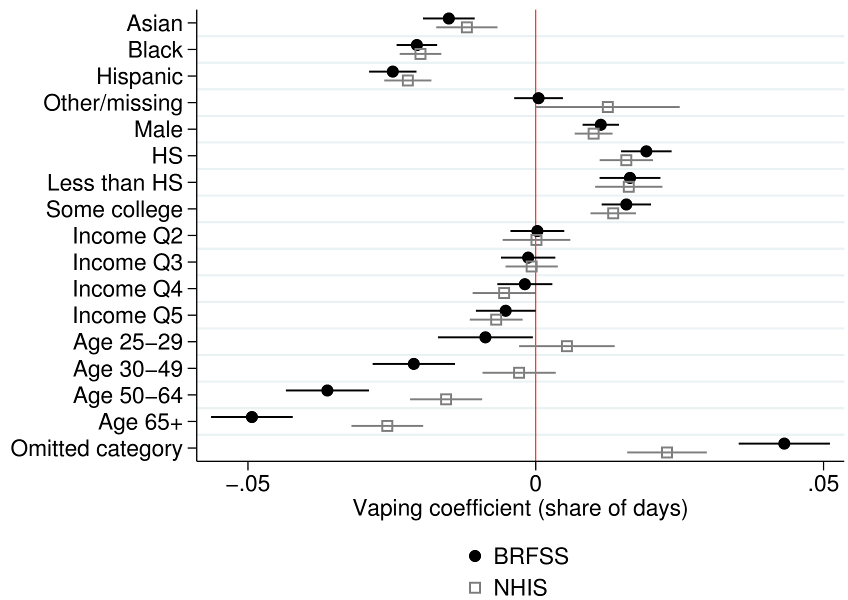

(b) **Youth**

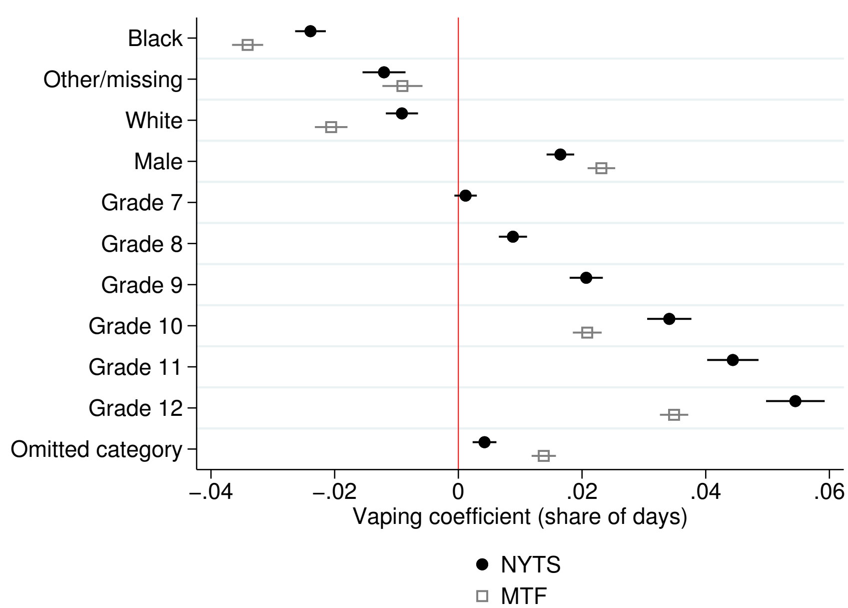

Notes: These figures present coefficients from Equation (20), a regression of vaping on demographic indicators, estimated separately by dataset. For adults, the omitted categories are White, female, college graduate, the lowest income quintile, and age group 18-24. For youth, the omitted categories are 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**

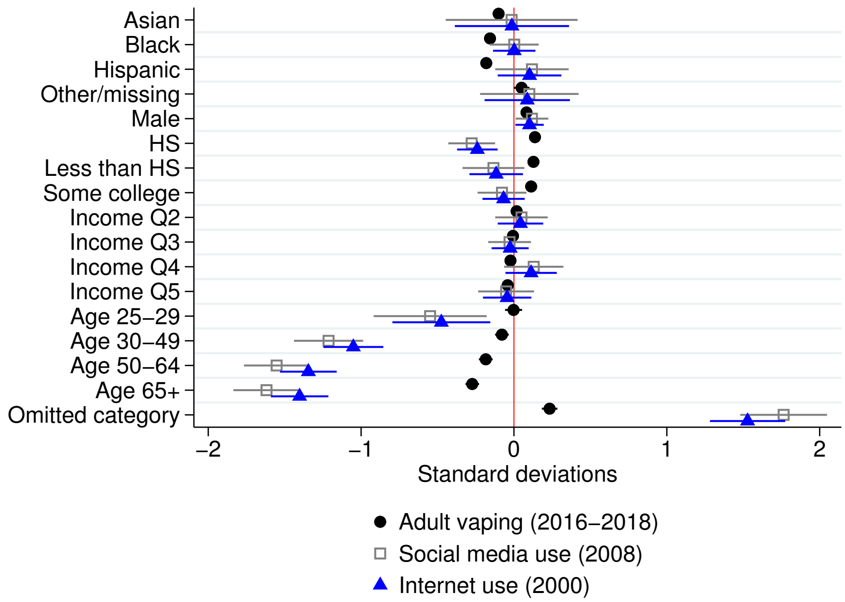

Notes: These figures present coefficients from regressions of vaping, social media use, or internet use on demographic indicators. Each dependent variable is normalized into standard deviation units for comparability. For adults, the omitted categories are White, female, college graduate, the lowest 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**

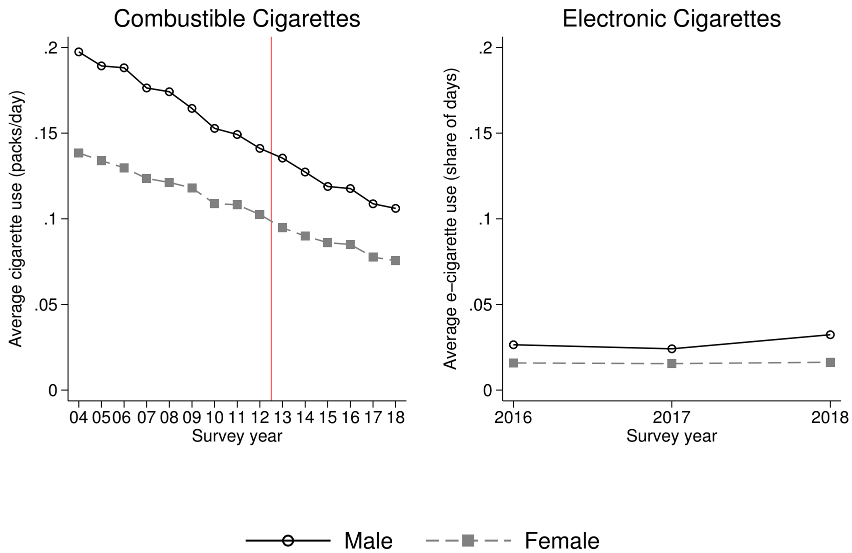

(b) **Youth**

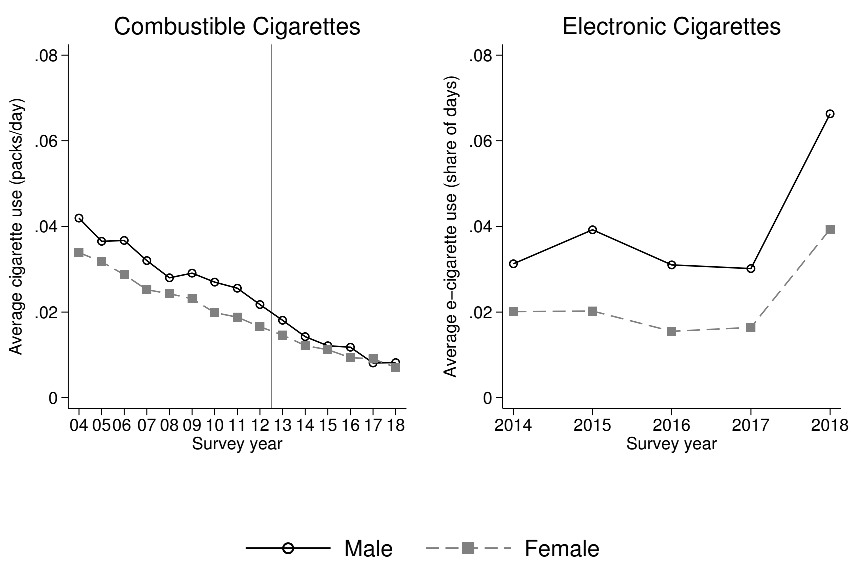

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

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

(a) **Adults**

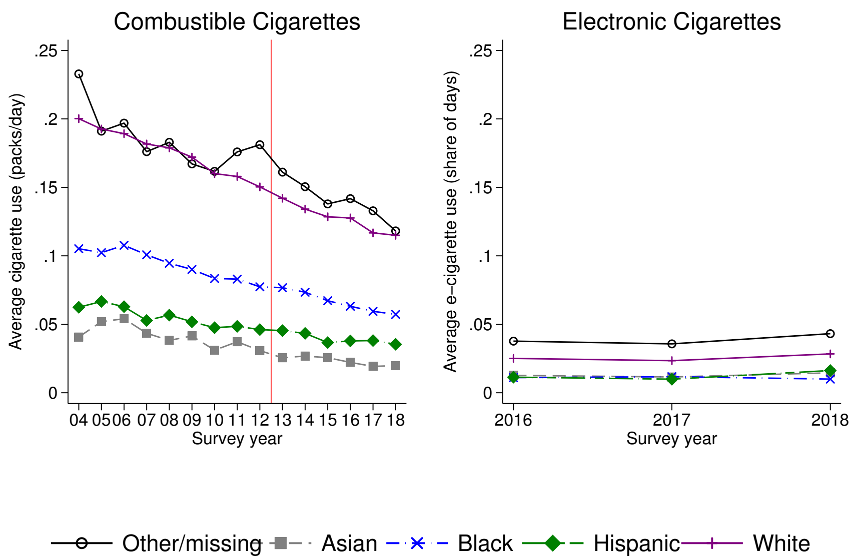

(b) **Youth**

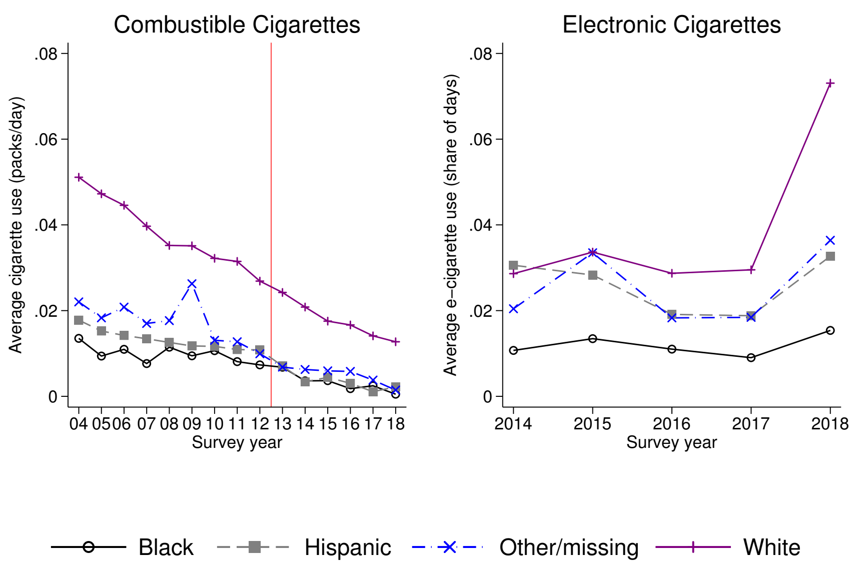

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

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

(a) **Adults**

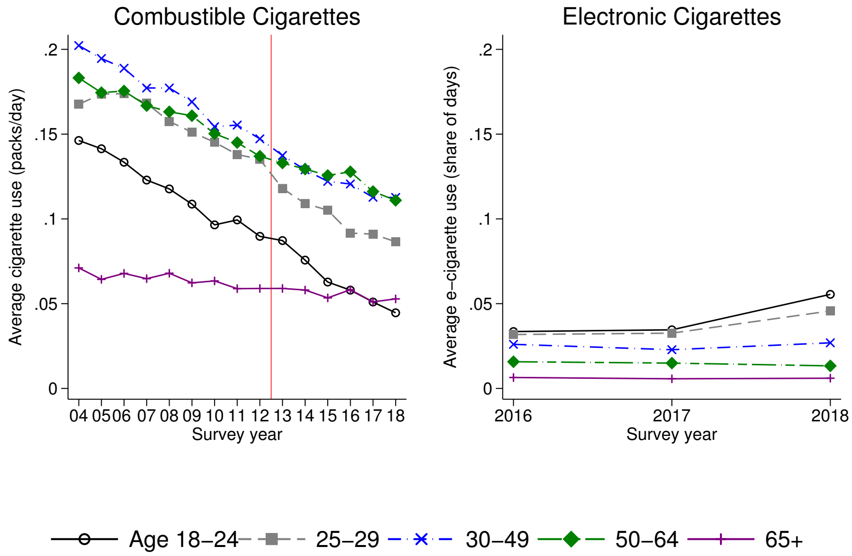

(b) **Youth**

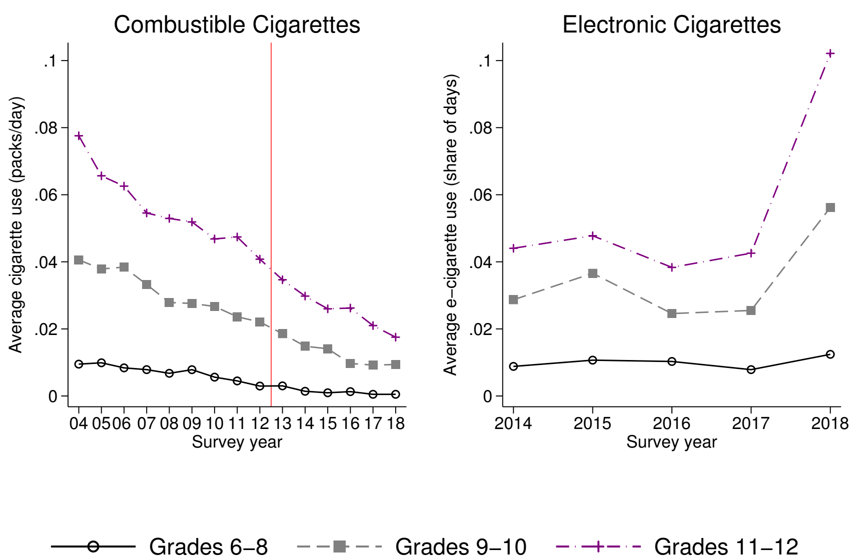

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

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

(a) **Education**

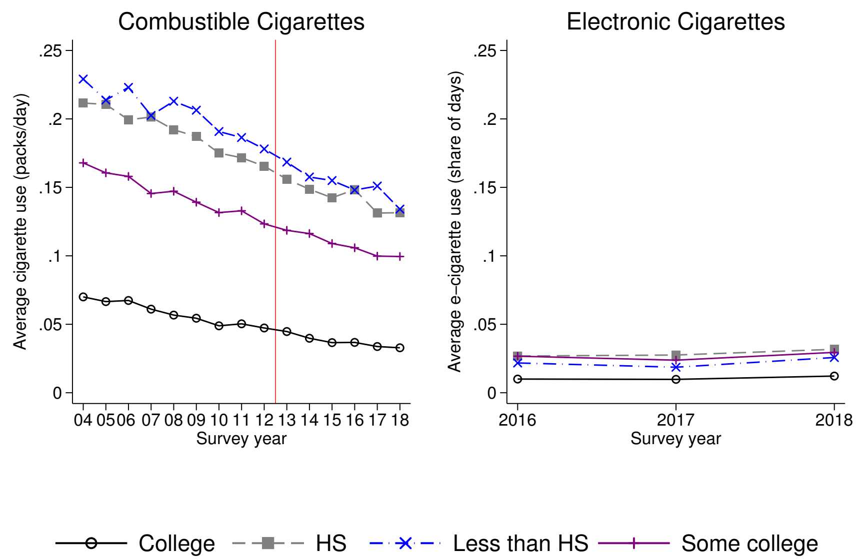

(b) **Income**

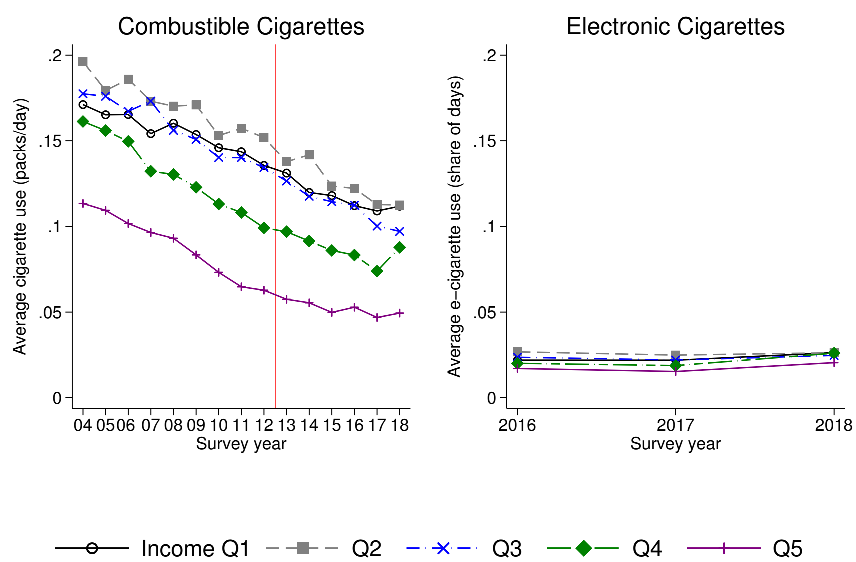

Notes: These figures present combustible cigarette and e-cigarette use by demographic group. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH.

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

(a) **Adults**

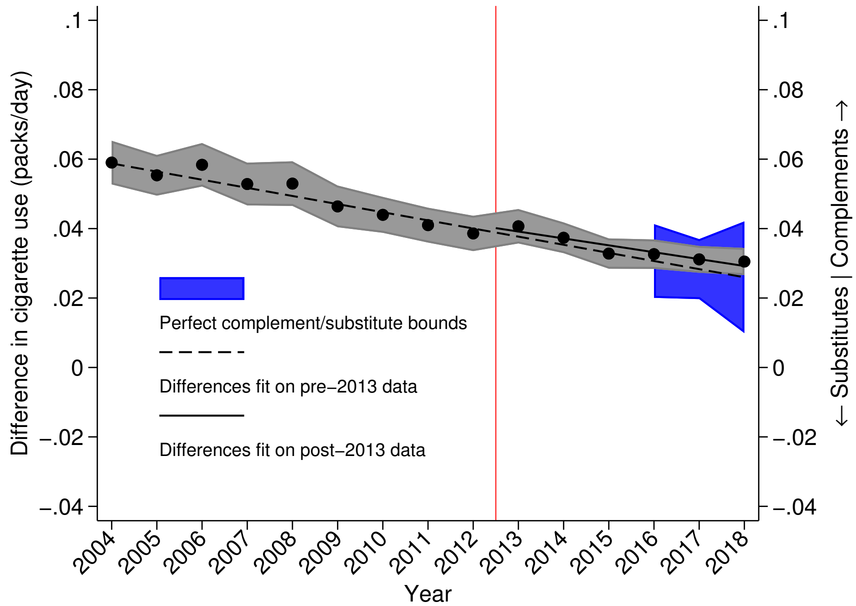

(b) **Youth**

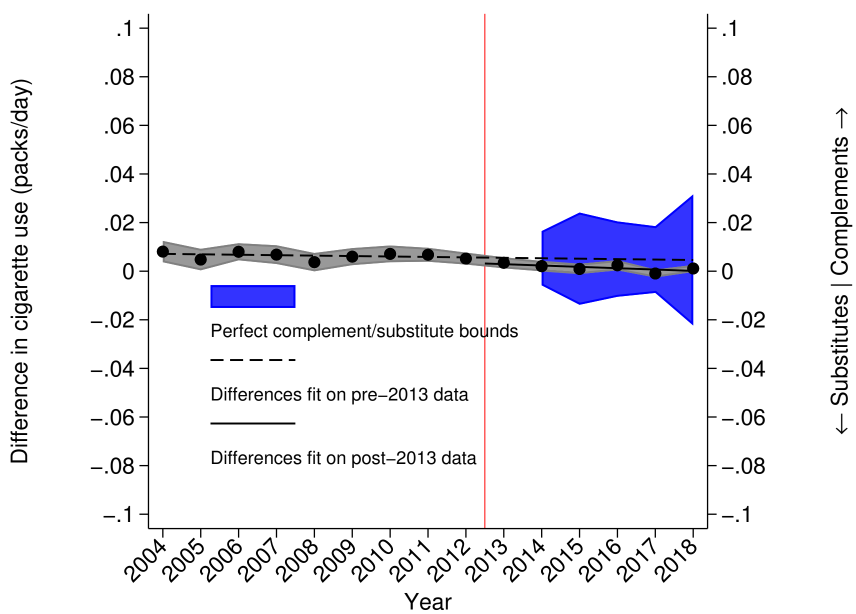

Notes: These figures present the difference in cigarette use for men versus women. Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette use to levels in the NSDUH. The perfect complement (substitute) bounds show 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**

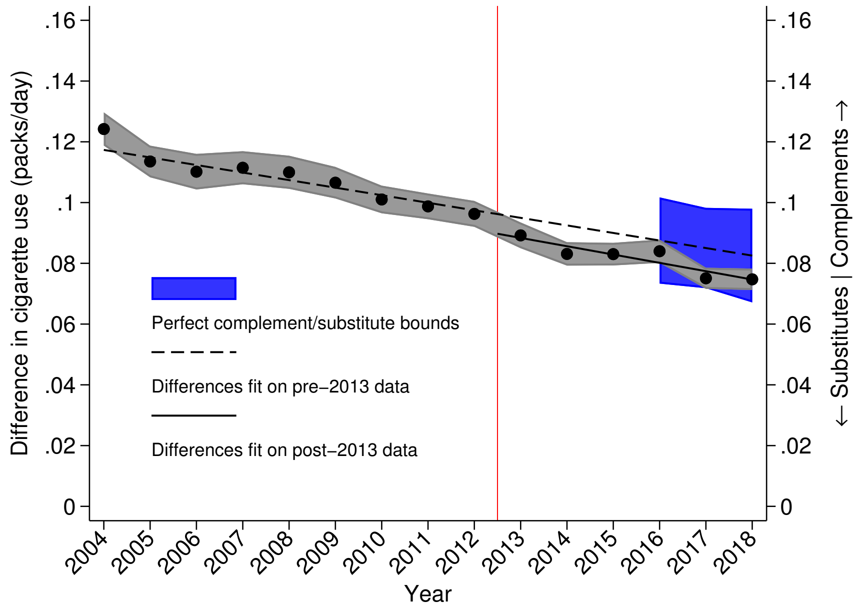

(b) **Youth**

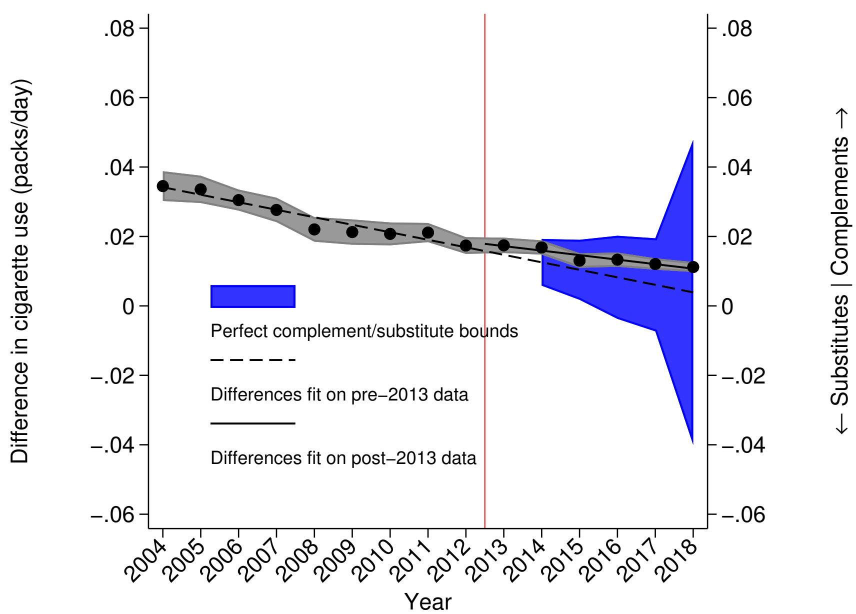

Notes: These figures present the difference in cigarette use for Whites and other races versus non-Whites (for adults) and Whites versus non-Whites (for youth). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette 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**

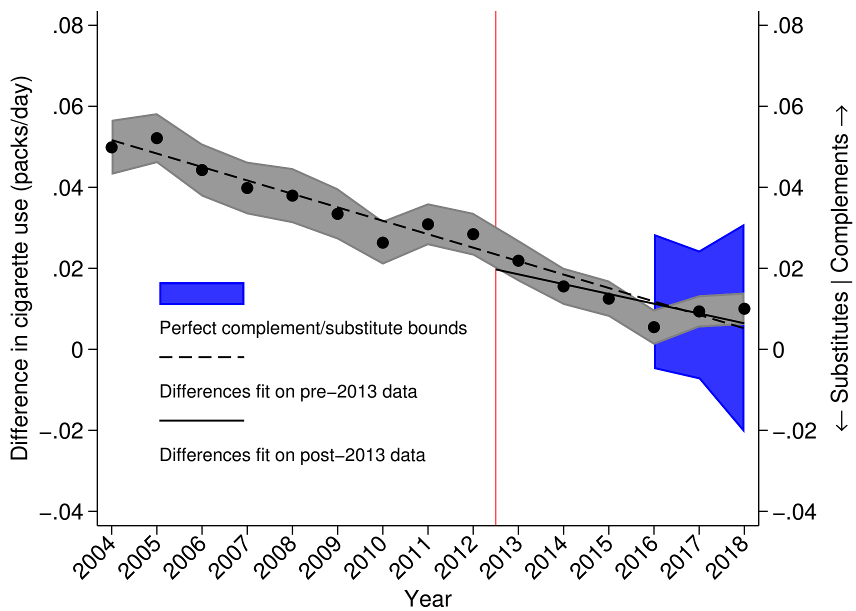

(b) **Youth**

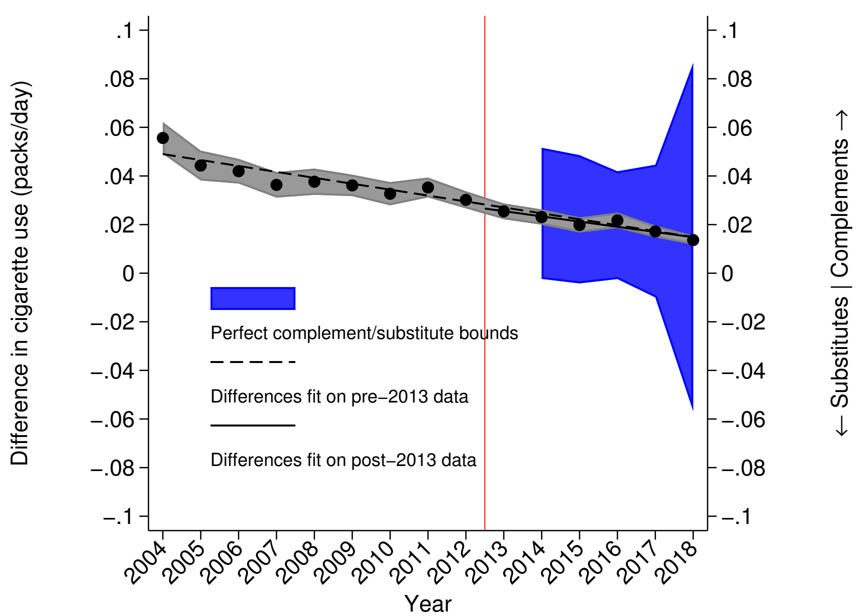

Notes: These figures present the difference in cigarette use by year for age ≤ 49 versus age ≥ 50 (for adults) and for grades ≥ 11 versus grades ≤ 10 (for youth). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS sampling frame change and rescale cigarette 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**

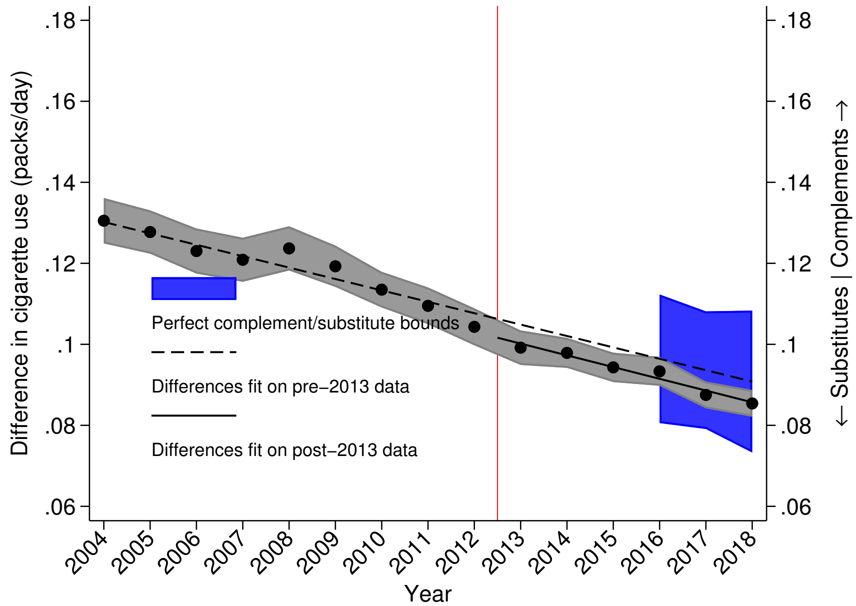

(b) **Income**

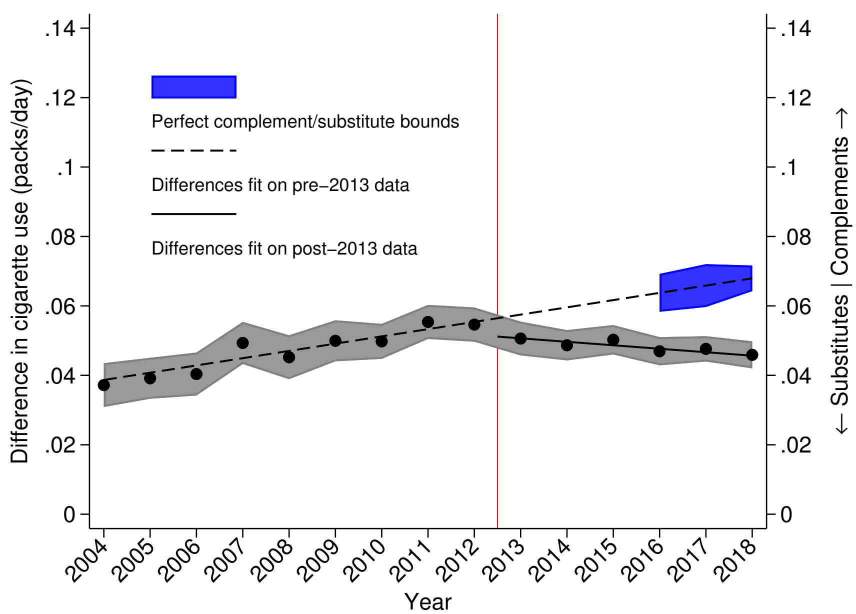

Notes: These figures present the difference in cigarette use by year for adults without versus with college degrees (Panel (a)) and adults in the bottom three versus top two income quintiles (Panel (b)). Average cigarette use for each group is residual of dataset controls, which address the 2011 BRFSS 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

$$\sigma_{\theta 2} := \frac{\chi^c}{\eta} \frac{q_\theta^c}{q_\theta^e} \tag{37}$$

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 $\hat{q}_\theta^j$, $\hat{\tilde{p}}^j$ 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:

$$s_{12} := Cov\left(\frac{\chi^c}{\eta}\frac{q_\theta^c}{q_\theta^e}, \frac{\chi^e}{\eta}\frac{\tilde{p}^e}{\tilde{p}^c}\Gamma\right) \tag{40}$$

$$= \chi^c \frac{q_\theta^c}{q_\theta^e}\frac{\tilde{p}^e}{\tilde{p}^c}\Gamma Cov\left(\frac{1}{\eta}, \frac{\chi^e}{\eta}\right) \tag{41}$$

$$\approx \chi^e \chi^c \frac{q_\theta^c}{q_\theta^e}\frac{\tilde{p}^e}{\tilde{p}^c}\Gamma V\left(\frac{1}{\eta}\right) \tag{42}$$

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**

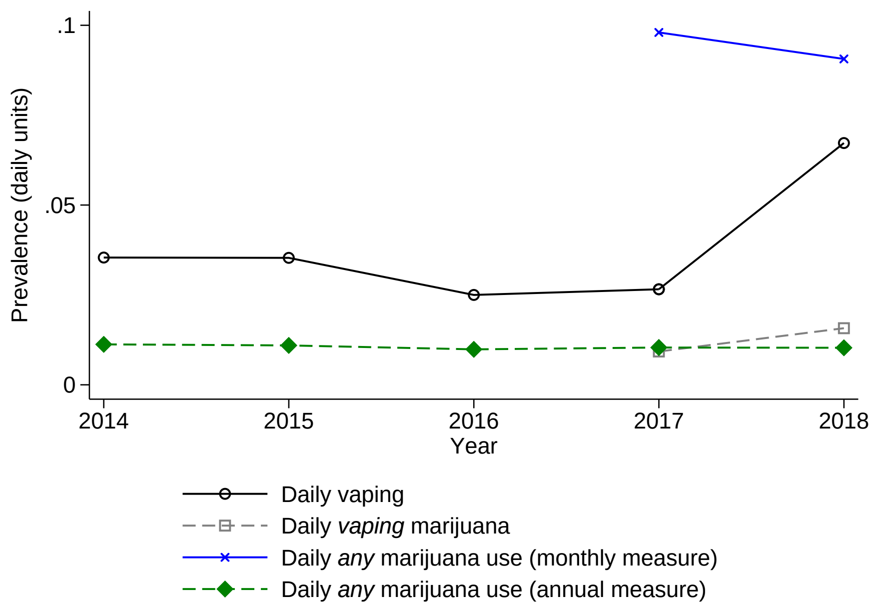

(b) **Grades 11–12**

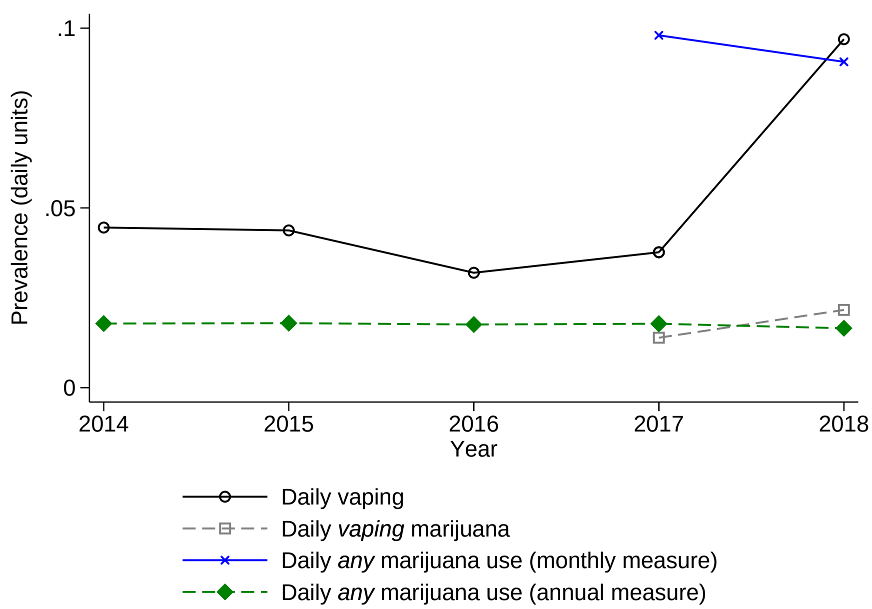

Notes: This figure presents trends in marijuana and e-cigarette use in the Monitoring the Future (MTF) survey. Panel (a) presents the full sample, while panel (b) focuses on grades 11 and 12. The black lines present our daily vaping measure. The gray lines present the average daily *vaping* marijuana 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**

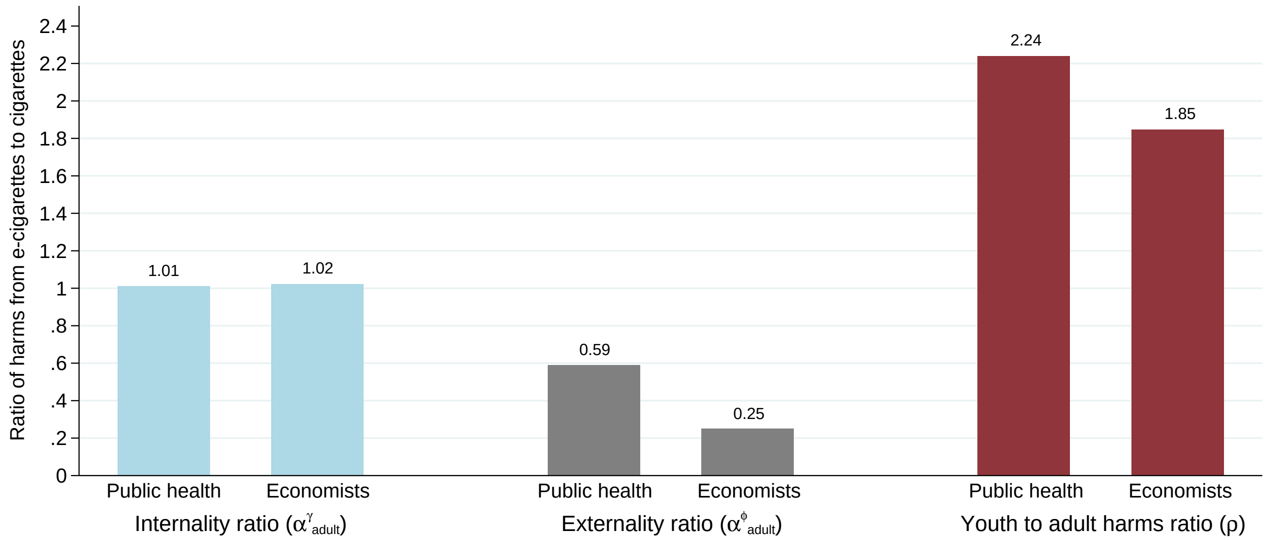

Notes: In our expert survey, we elicited the ratios of internalities and externalities from vaping relative to smoking and the ratio of uninternalized harms from vaping for youth relative to adults. This figure presents the averages of those ratios separately for public health researchers cited in National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

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

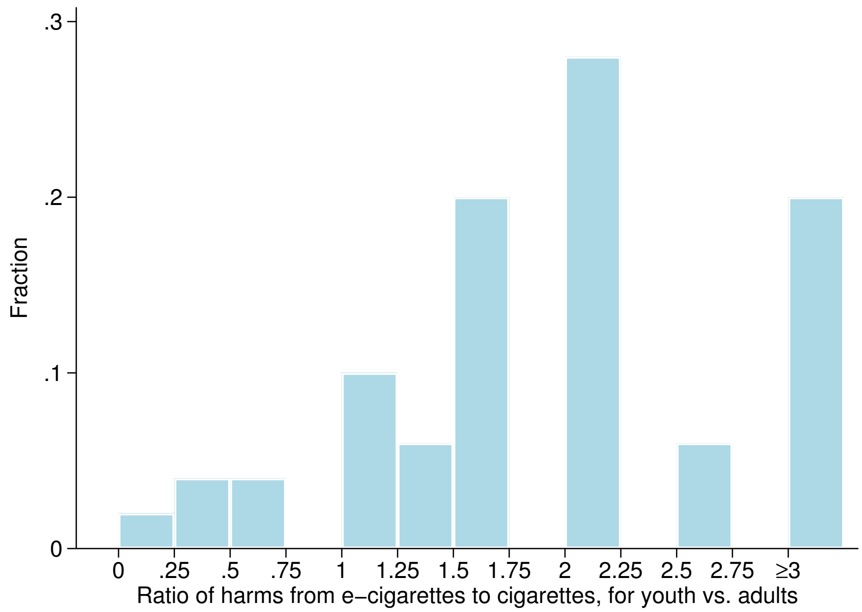

Notes: In our expert survey, we elicited the ratio of uninternalized harms from vaping for youth relative to adults. This figure presents the distributions of that ratio across experts.

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

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

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

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

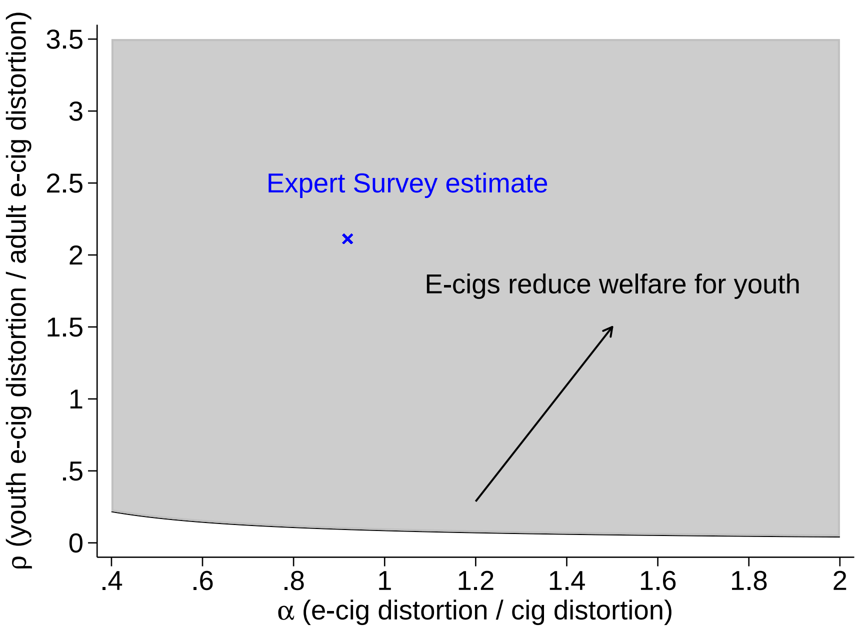

Notes: This figure presents parameter regions where a youth e-cigarette ban increases welfare, using Equation (15). All parameters other than α_{adult} and ρ are set at their means presented in Table 4.