Online Appendix:
Discrete Prices and the Incidence and Efficiency of Excise Taxes

Christopher T. Conlon and Nirupama L. Rao

January 23, 2020

A. Summary Statistics

We summarize our main (monthly) dataset in Table A1. For each state and product size (750mL, 1L, 1.75L) we report the number of store-product-month observations, the total sales, and the average price paid (total revenue divided by total sales). We exclude 1L bottles from Illinois and Louisiana because we have fewer than 8,000 such observations and they represent a very small fraction of sales; we keep them for Connecticut where they represent around 8% of the market. Additionally, we report the size of the tax increase for each state and product size which ranges from \$0.105 per bottle for 750mL bottles in Louisiana to \$1.87 per bottle for 1.75L bottles in Illinois. We also consider a weighted version of the same sample where we weight products by their annual sales in the same store for the calendar year prior to tax change. We use these weights because price changes are more important for more popular products.[1] The weighted sample has substantially lower prices because mass-market products are cheaper and more popular than high-end niche products.

Table A1: Summary Statistics by State and Size

			Unweighted Sample			Weighted Sample		
State	Size(mL)	$\Delta\tau$	# Obs	Total Sales	Price	# Obs	Total Sales	Price
CT	750	0.178	416,587	2,161,852	25.64	277,999	1,562,499	22.88
IL	750	0.802	6,547,716	48,319,333	18.53	3,464,535	31,259,575	16.05
LA	750	0.105	2,234,366	12,371,082	19.04	1,729,621	10,527,218	17.55
CT	1000	0.238	54,803	433,746	22.35	44,648	360,031	21.59
CT	1750	0.416	244,975	2,849,310	27.70	187,917	2,425,462	24.94
IL	1750	1.872	2,166,896	29,223,908	23.49	1,540,224	24,930,700	20.11
LA	1750	0.245	1,083,667	8,833,187	24.35	887,033	7,775,484	22.29

Note: Observations are store-month-UPC. Weights correspond to annual sales for the UPC and store during the calendar year prior to each state's tax change.

We observe that prices are broadly similar in Louisiana and Illinois, but substantially higher in Connecticut.[2]

[1] In total we observe 6,785 products many of which have extremely low sales. We consider restricting the sample to the top 1000 or top 500 products and it has almost no effect on any of the estimates we report.

[2] We examine the laws which facilitate collusive wholesale pricing in Conlon and Rao (2019).

B. Additional Heterogeneity in Pass-Through Rates

The pass-through rates detailed in Table 2 average over all products in each state of the specified size. These averages belie some sources of meaningful heterogeneity in pass-through rates.

Table B1 details the substantial heterogeneity across stores in their response to the tax. Stores that sell products at relatively lower prices (in the prior month) are more likely to raise prices in response to the tax change than those stores with relatively higher prices. We use two different discrete measures of high and low-price stores by product and month: the first column uses dummies for prices above or below the median price, the second column uses dummies for the highest and lowest price retailer (allowing for ties) selling the same product. For both measures we find that at low-price stores the tax is passed on at a rate of roughly 260% to 270% while at high-price stores pass-through point estimates are below 10% and statistically indistinguishable from zero. We also employ a continuous measure of relative price in the third column and again find lower relative prices are correlated with larger pass-through rates.

Table B1: Pass-Through: Taxes to Retail Prices Relative to Other Stores

	Above/Below Median	Min/Max	Continuous
Δ Tax	1.045*	1.132*	1.128*
	(0.270)	(0.237)	(0.221)
Δ Tax * High	−0.923*	−1.141*	
	(0.352)	(0.366)	
Δ Tax * Low	1.548*	1.584*	
	(0.411)	(0.489)	
Δ Tax * Relative			−0.177*
			(0.031)
High Price	−0.381*	−0.350*	
	(0.050)	(0.042)	
Low Price	0.152*	0.246*	
	(0.029)	(0.031)	
Relative to Median			−0.034*
			(0.003)
Observations	427,957	427,957	427,957
Adjusted R^2	0.028	0.025	0.031
Product FE	Yes	Yes	Yes
Month+Year FE	Yes	Yes	Yes
High Measure	Above Median	Maximum	Continuous
Low Measure	Below Median	Minimum	% Deviation

* Significant at the 1 percent level

Note: The table above reports OLS estimates of the pass-through of taxes into retail prices in Connecticut for low and high-priced stores. All regressions weighted by 2011 Nielsen units (normalized by size) and reported standard errors are clustered at the UPC level. Relative prices are from the *prior* month. Columns (1)+(2) use indicator variables. Column (3) uses percentage deviation from median price.

Table B2 examines how retail price changes are modulated by the cumulative change in the lowest wholesale price of a product since a store's last retail price change. This state variable captures the *pressure* to adjust prices that a retailer faces from the build-up of wholesale price changes. In column 1 of Table B2 we include the wholesale state variable alongside the size-interacted tax change variable; in column 2 we include an interaction between the wholesale state variable and the change in tax and in column 3 we include both the main and interaction terms. In all specifications larger cumulative changes in wholesale prices since the last retail price change lead to larger retail price increases. While the main effect and interaction term are statistically significant on their own (columns 1 and 2) when both are included only the main effect of the wholesale state variable is significant though the point estimate of the interaction term is of the expected positive sign.

Table B2: Pass-Through: Adding Wholesale Price Change State Variable (CT only)

	(1)	(2)	(3)
Δ Tax (750mL)	2.840*	3.147*	2.886*
	(0.772)	(0.801)	(0.785)
Δ Tax (1000mL)	1.868*	2.013*	1.909*
	(0.527)	(0.551)	(0.534)
Δ Tax (1750mL)	0.522	0.489	0.598
	(0.413)	(0.419)	(0.416)
ΔW	0.113*		0.114*
	(0.016)		(0.016)
$\Delta W \times \Delta T$		0.258*	−0.056
		(0.084)	(0.086)
Observations	95,136	95,136	95,136
Adjusted R^2	0.060	0.032	0.060

* Significant at the 1 percent level.
Note: The table above reports quarterly OLS estimates of the pass-through of taxes into retail prices in Connecticut controlling for the change in the product's wholesale price since the last change in the store's retail price for that product. All regressions weighted by 2010 Nielsen units (normalized by size) and reported standard errors are clustered at the UPC level.

C. Data

C.1. Aggregation of Nielsen Weekly Data to Monthly Data

The Nielsen scanner data are recorded weekly, and some weeks span two months. We aggregate the data to the monthly level for the initial analysis for a number of reasons. The first is that in Connecticut, wholesale (and retail) prices are not allowed to vary within a month. This is not necessarily true in Illinois or Louisiana where prices can adjust more flexibly. Second, when tax changes are observed, the occur on the first day of the month. We allocate weeks to months based on the calendar month of the last day of the corresponding week. When we aggregate, we take the last price (Nielsen revenue divided by units) recorded in each month and total sales for each product-store-month. In practice, there is only a single price for 99% of store-month-product observations in Connecticut once we exclude the first week of the month (which may contain data from two months).

C.2. Consolidation

We consolidate products so that a product is defined as brand-flavor-size such as *Smirnoff Orange Vodka 750mL*. Sometimes a "product" may aggregate over several UPC's, as changes in packaging can result in a new UPC. UPC changes most commonly arise with special promotional packaging such as a commemorative bottle, or a holiday gift set. At other times, the change in UPC may be purely temporal in nature. A product may also be available in both glass and plastic bottles at the same time. We rarely observe price differences for glass and plastic packaging within a product-month, so we also consolidate these UPCs.

In total, these consolidations help us to construct a more balanced panel of products over time, and avoid gaps during holiday periods, or products going missing when packaging changes. This is especially important when our goal is to capture *changes* in prices within a product-store over time.

C.3. Cleaning Prices

Nielsen data report weekly sales at the store-UPC level. Prices are not observed directly but rather imputed from revenues as $p_t \approx \frac{p_t \cdot q_t}{q_t}$. It is common to adjust or filter prices under a number of scenarios: (1) transitional prices (2) temporary sales (3) clearance/closeouts. For (1) observed prices may not represent transaction prices, but rather the weighted average of two different price points. For (2) and (3) the observed prices may in fact be transaction prices, but those prices may not end in 0.99 as many stores/chains use unusual price endings internally to track sales or clearances.

It is helpful to consider a sequence of prices for a product $[p_{t-2}, p_{t-1}, p_t, p_{t+1}, p_{t+1}]$ and so on. In many cases p_t and p_{t+1} are not adjacent weeks but rather adjacent periods in which a sale is recorded. This is is important because Nielsen does not record any information unless a product was purchased that week.

Rule #1: Transitional Prices

If the store changes its prices at the end of a Nielsen-week (Saturday to Saturday) the the recorded price should match the actual transaction price. A more likely scenario is that a store changes its price midweek so that revenues include sales recorded at p_{t-1} and p_{t+1}, while no transactions actually take place at the recorded p_t. Because we are interested in *price changes* we will replace p_t with the closer of p_{t-1}, p_{t+1}, so that our recorded price corresponds to an actual transacted price (rather than a weighted average).

We can detect these transitional prices by when $p_t \in (p_{t-1}, p_{t+1},)$ or $p_t \in (p_{t+1}, p_{t-1},)$ and then we use observed sales q_t to construct a convex set of potential prices where $w = [1, \ldots, q_t - 1]/q_t$ and $p_t \approx w p_{t-1} + (1-w) p_{t+1}$ to see if p_t lies on the grid of transitional prices.

We apply this rule only to price endings not in the four most commonly used price endings for each chain in our dataset.

We include the code below:

```
def transition_prices(p1,p2,pobs,q):
    # is observed price between (p1,p2)
    if not (((pobs < p1) & (pobs > p2)) | ((pobs > p1) & (pobs < p2))):
        return np.nan
    #convex weights
    w=np.array([x/q for x in range(0,q+1)])
    # grid of possible prices
    possible=w*p1 + (1-w)*p2
    # check if observed price is within 1 cent of rounded grid of possible prices
    # then return the closer price --> otherwise missing
    if any(np.isclose(pobs,np.round(possible,3),atol=.01)):
        if np.abs(p1-pobs) <= np.abs(p2-pobs):
            return p1
        else:
```

```
        return p2
    else:
        return np.nan
```

Rule #2: Temporary Sales

There is evidence that many stores (or chains) use prices with unusual endings as an internal way to track temporary sales or promotional prices. For example, *reference prices* may end in 0.99 or 0.49 but temporary sales may end in 0.97 or 0.12. This is largely not an issue in Connecticut where regulations make temporary sales relatively rare, but appears to be more common in Illinois and Louisiana where retailers are free to increase and reduce prices at will.

We consider a price a one-week temporary sale if: $p_{t-1} = p_{t+1} = p_{t+2} = p_{t-2}$ and $p_t < p_{t-1}$. For a two week temporary sale, we define the first week as: $p_{t-3} = p_{t-2} = p_{t-1} = p_{t+2} = p_{t+3}$ and $p_t < p_{t-1}$. For both cases we require that the final observed price for a product cannot be considered a "sale" price.

```
def add_temp_sales(df):
   # observed leads and lags are all identical except p_t, p_{t+1} and p_t < p_{t-1}
    x=(df[['p_lag1','p_lag2','p_lag3','p_lead2','p_lead3','p_lead4']].std(axis=1)==0) & ...
     (df['price']< df['p_lag1']) & ~(df['p_lead2'].isnull())
   # p_t == p_{t+1}
    x=x & (df['price'] == df['p_lead1'])

    # Detect the (optional) second week of sale
    y=(df['p_lag1'] == df['price']) & x.shift(1)
    # Detect one week sales
    z=(df[['p_lag1','p_lead1','p_lag2','p_lead2']].std(axis=1)==0) ....
     &(df.price<df.p_lag1)&(~df.p_lead1.isnull())
    df['sales_2wk']=(x|y)
    df['sales_1wk']=z
    return df
```

Rule #3: Closeouts/Clearance Items

We also find and tag clearance or closeout prices. These are the final price at which a good transacts p_T subject to some conditions. We look for cases where the last price at which a good transacts is otherwise unobserved in the full history of prices $[p_1, p_2, \ldots, p_T]$ except in a sequential run of prices at the end of the dataset. Thus it must be that $p_{t-k} = p_T$ for all $k = 0, 1, \ldots, K$ and some $K \geq 1$. It must also be that $p_{t-K} < p_{t-K-1}$ (the first price beginning the run of clearance prices is lower than the previous price).

This rule is conservative in that if a product starts at 19.99 and is reduced to 14.12 and then later to 12.15 only the 12.15 price is considered a closeout, and 14.15 is also not considered a temporary sale (since the price does not return to 19.99).

```
# return vector same size as p
def np_closeout(p):
    p=p.values
    y=np.zeros(len(p))

    idx=np.where(p==p[-1])
```

```
run_start=idx[0][0]
#contiguous set of prices that match the final price (without gaps)
if (np.diff(idx)==1).all() & (p[run_start] < p[run_start-1]):
    y[idx]=1
return y
```

C.4. Wholesale Data

We draw on a hand-collected dataset of wholesale prices for the state of Connecticut. Wholesale prices are a key predictor of retail price changes. These prices were scraped by us from the Connecticut Department of Consumer Protection (DCP) from August 2007 to August 2013. These data are available because Connecticut requires that all licensed wholesalers post prices. Wholesalers agree to charge retailers these prices for the entire month, and are legally not allowed to provide quantity discounts or price discriminate.[3] Only 18 wholesale firms have ever sold brands of distilled spirits that we observe in the Nielsen dataset, and more than 80% of sales come from just six major wholesalers. Because Illinois and Louisiana do not require that wholesalers publicly post prices, we do not have wholesale pricing information for these states.

For our welfare calculations we use estimates from Conlon and Rao (2019) of the prices paid by wholesale firms to importers and distillers. These marginal costs are estimated at the product level using a structural model of demand in the approach of Berry et al. (2004).

D. Quarterly Pass-Through and Price Points

D.1. Constructing Quarterly Data

We estimate our nonlinear models using quarterly data to avoid the repetitive use of monthly observations. We allocate weeks to quarters using the last day of each week. This works for Connecticut where the tax changes on July 1 (the first day of Q3), and Louisiana where the tax changes on April 1 (the first day of Q2). We have to modify this procedure for Illinois, where the tax changes on September 1. For Illinois, we change the starting month for each quarter so that Q1 begins in December, Q2 begins in March, Q3 begins in June, and Q4 begins in September. This way, the tax change happens at the beginning of of Q4 under our adjusted definition.[4] When we aggregate, we take the last price (Nielsen revenue divided by units) recorded in each quarter and total sales for each product-store-quarter.

For comparison, in the left panel of Table D1 we reproduce Table 2 from the main text using the quarterly data instead of monthly data. We find that the patterns are broadly similar, though we have fewer observations, larger standard errors, and less precise control over seasonality. In all but once case (1750mL bottles in Connecticut) we find evidence that taxes are over-shifted $\rho > 1$ and that conditional on a price change, estimated pass-through rates are higher. The latter effects are muted as one might expect because at the quarterly level price changes are more common overall.

[3] Connecticut is one of 12 states with a set of regulations known as *Post and Hold*, which mandates that all wholesalers post the prices they plan to charge retailers for the following month. Wholesalers must commit to charging these prices for the entire month (after a look-back period when wholesalers can view one another's initially posted prices and adjust their prices downwards without beating the lowest price for the product). For a detailed analysis of these regulations please see Conlon and Rao (2019).

[4] These definitions make it difficult to include quarterly fixed effects in regression specifications, as December (a high sales month) is in Q4 for (CT,LA) and Q1 in IL. For this reason we only ever consider $State \times Quarter$ fixed effects.

Table D1: Pass-Through: Taxes to Retail Prices (Quarterly)

	Quarterly		w/ Price Points	
	All Observations	Δ Retail Price ≠ 0	All Observations	Δ Retail Price ≠ 0
Connecticut July 1, 2011 Tax Increase of $0.24/L				
Δ Tax (750mL)	2.944*	4.476*	2.723*	4.031*
	(0.735)	(1.716)	(0.532)	(1.342)
Δ Tax (1000mL)	2.094*	2.639	1.952*	2.568
	(0.509)	(1.227)	(0.437)	(1.130)
Δ Tax (1750mL)	0.800	0.950	0.766	0.929
	(0.373)	(0.774)	(0.359)	(0.765)
Illinois Sept 1, 2009 Tax Increase of $1.07/L				
Δ Tax (750mL)	2.738*	3.447*	2.385*	3.130*
	(0.224)	(0.256)	(0.190)	(0.218)
Δ Tax (1750mL)	1.375*	1.645*	1.226*	1.470*
	(0.099)	(0.123)	(0.096)	(0.115)
Louisiana April 1, 2016 Tax Increase of $0.14/L				
Δ Tax (750mL)	1.791	3.935	1.983	4.434*
	(1.112)	(1.759)	(1.002)	(1.696)
Δ Tax (1750mL)	1.198	2.159	1.632*	3.004*
	(0.576)	(0.906)	(0.445)	(0.735)
Observations	3,035,603	1,606,649	2,948,414	1,443,105
Adjusted R^2	0.052	0.094	0.074	0.141

* Significant at the 1% level.
Note: We include fixed effects for UPC, quarter of year, and year (all interacted with state). All regressions are weighted by annual sales for the UPC and store for the year prior to the tax change. Standard errors are clustered at the state-UPC level. Price points are defined in Table D3.

D.2. Constructing Price Points

When we estimate the ordered logit models, we consolidate several price changes into larger bins. We report the data in relatively narrow intervals in Table D2. Some important patterns emerge. First, the majority of price changes are within a few cents of zero. Second, as we have documented in the main text, the most popular price change intervals are the ones that contain $-\$1.00, +\$1.00, +\$2.00$. There are some other important patterns that are worth mentioning. (1) Connecticut price changes are more likely to be in or around whole dollar increments than price changes in other states. As described in the main text, this is in part because Connecticut regulations prevent mid-month price changes that reduce our ability to accurately measure prices from the revenue and quantity information reported by Nielsen. In addition, Connecticut limits temporary sales which mean that most reported prices coincide with transaction prices. (2) There are a relatively small fraction of very large price changes (both positive and negative) which we will ultimately ignore when estimating our ordered logit models. These may represent undetected closeouts, or substantial departures from previous pricing strategy and often affect low sales high price specialty products (with prices $\geq \$150$). Other non whole-dollar price changes such as 50 cents, represent less than 2% of the observations in the overall data. Thus aggregating over these price changes is relatively innocuous. (3) moderate positive price changes of \$3.00 or more are relatively uncommon except in Illinois during the quarter of the tax change.

We further consolidate the data in Table D2 to a set of price change increments which we use in our ordered logit estimation. As before, we assign each price change to an interval, and report its frequency in the table. We also can assign each interval a value or "Category Value", these values

Table D2: Frequency of Quarterly Price Change Intervals

	All Quarters			Quarter of Tax Change		
Price Change	CT	IL	LA	CT	IL	LA
(-10,-6.03]	0.51	0.85	0.57	0.30	0.82	0.74
(-6.03,-5.97]	0.18	0.36	0.31	0.11	0.24	0.44
(-5.97,-5.03]	0.092	0.25	0.15	0.12	0.16	0.33
(-5.03,-4.97]	0.31	0.60	0.64	0.44	0.30	0.77
(-4.97,-4.03]	0.11	0.48	0.25	0.03	0.33	0.38
(-4.03,-3.97]	0.57	1.11	1.26	0.62	0.90	1.67
(-3.97,-3.03]	0.19	0.81	0.49	0.21	0.54	0.57
(-3.03,-2.97]	0.86	2.25	2.14	0.84	1.48	2.26
(-2.97,-2.03]	0.33	1.37	1.00	0.20	0.85	1.47
(-2.03,-1.97]	1.50	4.84	3.95	1.63	0.94	4.01
(-1.97,-1.03]	0.44	2.29	1.86	0.44	1.16	2.77
(-1.03,-0.97]	1.85	8.03	6.70	4.22	3.63	7.93
(-0.97,-0.53]	0.15	1.24	0.93	0.14	0.67	0.90
(-0.53,-0.47]	0.10	1.08	1.23	0.14	0.56	1.28
(-0.47,-0.03]	0.28	1.57	1.22	0.45	0.91	1.38
(-0.03,0.03]	77.45	43.55	50.73	58.18	14.60	38.77
(0.03,0.47]	0.71	1.64	1.60	0.81	1.54	1.34
(0.47,0.53]	0.63	1.30	1.92	1.30	1.23	3.72
(0.53,0.97]	0.41	1.31	0.97	1.48	1.72	1.11
(0.97,1.03]	5.88	9.32	8.28	14.85	13.46	11.93
(1.03,1.97]	0.62	2.42	2.18	1.42	6.52	3.06
(1.97,2.03]	2.35	4.92	4.29	4.46	16.25	4.95
(2.03,2.97]	0.36	1.43	1.10	0.89	4.92	0.99
(2.97,3.03]	1.10	2.36	2.22	1.87	9.66	2.69
(3.03,3.97]	0.23	0.70	0.46	0.48	2.45	0.47
(3.97,4.03]	0.72	1.13	1.07	1.40	4.07	1.27
(4.03,4.97]	0.15	0.40	0.23	0.54	1.35	0.42
(4.97,5.03]	0.42	0.69	0.74	0.65	4.22	0.60
(5.03,5.97]	0.059	0.22	0.13	0.20	1.10	0.081
(5.97,6.03]	0.23	0.37	0.31	0.35	0.96	0.38
(6.03,6.97]	0.056	0.12	0.071	0.11	0.32	0.062
(6.97,7.03]	0.14	0.22	0.16	0.075	0.56	0.18
(7.03,10]	0.39	0.31	0.29	0.30	0.75	0.36
$\|\Delta P\| > 10$	0.60	0.46	0.53	0.75	0.81	0.70

are *not* used in the ordered logit estimation, however we do use these values when calculating our counterfactuals.

We estimate quarterly pass-through regressions using these transformed values and compare the results to the un-transformed quarterly data in Table D1. The right panel of Table D1 confirms that transforming the data this way has little impact on the estimated pass-through rates. In all cases pass-through estimates using the transformed data are statistically indistinguishable from the quarterly pass-through estimates using the un-transformed data.

We can examine the in-sample fit of the ordered logit by looking at Connecticut during the month of the tax change. This period is important because it is the period we use for all of our counterfactual experiments. In Figure D1 we compare the observed price changes to those predicted under our main ordered logit specification. In general, the fit of the model is good, though we under-predict zero price changes and over-predict both positive and negative price changes. We also see that that the fraction of observations at large price changes $> +\$3$ or $< -\$3$ is fairly small, which suggests restricting the domain of potential outcomes may not be a major problem. Likewise, we see that nearly all price changes adhere to the grid of pre-specified price points with the possible exception of +\$0.50, which we examine in more detail in a robustness test below.

Table D3: Frequency of Quarterly Price Changes with Price Points

		All Periods			During Tax Change		
Category Value	Interval	CT	IL	LA	CT	IL	LA
-2	(-5.1,-1.5]	4.11	12.78	10.90	4.19	6.10	12.74
-1	(-1.5,-0.25]	2.43	12.12	10.17	5.02	5.67	12.27
0	(-0.25,0.25]	77.97	45.23	52.32	58.68	16.04	40.03
$+\frac{1}{2}$	(0.25,0.75]	1.31	2.71	3.07	2.94	2.81	5.22
+1	(0.75,1.5]	6.45	11.59	10.38	16.27	18.29	14.52
+2	(1.5,2.5]	2.74	6.64	5.67	5.41	21.98	6.45
+3	(2.5,3.5]	1.40	3.29	2.81	2.35	12.86	3.28
+4	(3.5,4.5]	0.89	1.63	1.37	2.05	5.91	1.72
+5	(4.5,5.1]	0.47	0.86	0.82	0.78	4.65	0.67
	$\|\Delta P\| > 5.1$	2.23	3.14	2.50	2.29	5.69	3.11

We can also look at how incorporating additional price points affects our estimates of predicted price changes and pass-through in the ordered logit specification. We use the same quartic polynomial in $\Delta\tau$ as before, but now instead of restricting $\Delta p \in \{\leq -1, 0, 1, 2, \geq 3\}$, we allow for additional price points at $\Delta p \in \{+0.5\}$ in one specification and further adds $\Delta p \in \{-2, +4\}$ in another. We report the results for predicted price changes in Figure E5. We find that adding the additional price point at fifty cents has a negligible effect on our predicted price changes, and that adding the additional price points at $\{-2, +4\}$, leads to slightly lower predictions for small tax changes (because of the -2) and slightly higher predictions for large tax changes (because of the $+4$). The overall qualitative patterns are preserved. We prefer to consolidate the $+\$4, +\$5$ price changes with the $+\$3$ price change because outside of Illinois during the month of the tax change they are very rare and thus difficult to predict accurately.

E. Additional Robustness Tests

Here we consider additional robustness tests.

We vary the elasticity of demand used to calculate counterfactual welfare: deadweight loss, producer surplus, consumer surplus, tax revenue, and incidence. We replicate Figure 7 but instead

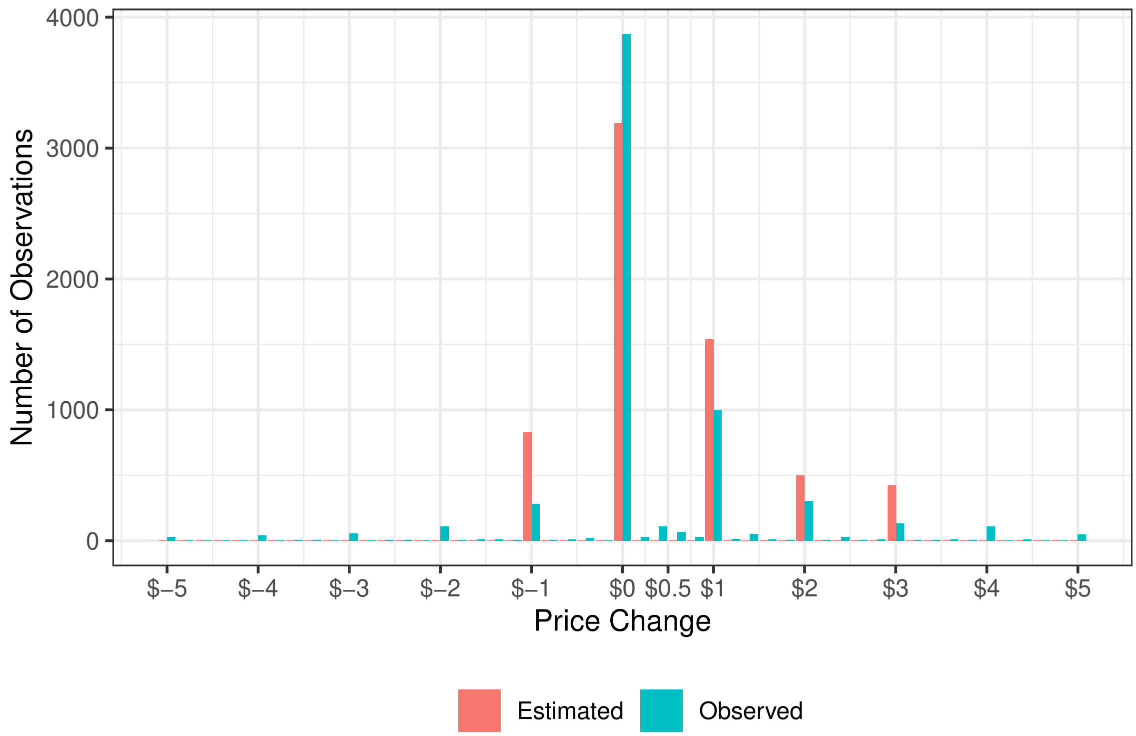

Figure D1: In-sample fit of Ordered Logit

Note : The figure above compares the predictions of our ordered logit model to the price changes observed in Connecticut during the month of the tax change. The predictions correspond to our preferred ordered logit model, which employs a quartic orthogonal polynomial of the tax change. The controls used in the ordered logit model measure the change in wholesale price since the last change in retail price, total sales by product overall stores, total sales by store over all products, the natural log of the price for the product the prior quarter at the same store, whether that store sold the product at the highest or lowest price the prior quarter and the difference between the price last quarter and the median price across all stores last quarter. The regression also includes state-varying controls; specifically, it includes state fixed effects and interactions between state dummies and Total Product Sales, Total Store Sales, log Lag Price, High Price, Low Price and the Relative Price cubic polynomial. Weights are balanced by state, bottle size and tax change indicator.

use elasticities $\epsilon_d \in \{-2.5, -4.5\}$. These are meant to capture the range of product-level own price elasticities reported in the empirical literature. See Conlon and Rao (2019) or Miravete et al. (2018). We provide those results in Figures E1, E2. As we might expect, as we increase the elasticity of demand, consumers bear less of the burden and firms bear more. The social cost of taxation responds the opposite way in that as demand becomes more elastic, ΔDWL increases. For the less elastic demand, the linear pass-through estimate lies strictly above the price points/ordered logit estimates, and for more elastic demand it lies (partly) slightly below. In both cases the qualitative patterns remain similar. Both the efficiency and incidence calculations produce a series of U-shaped curves as we increase the size of the tax that qualitatively match our main result.

We also vary the markups used in the welfare calculations to estimate MC_j. In the main text we assume $\mu = 1.5$, here we also consider $\mu = 1.2$ and $\mu = 2.0$ in Figures E3 and E4 respectively. Increasing the markup increases ΔPS at all values of the tax increase, and increases the firm portion of ΔDWL. Again, while the scale of the y-axis varies with the markup term, the qualitative predictions of our model remain the same.

The results are also insensitive to enlarging the set of potential price points. In Figure E5 we plot the change in price predicted by the baseline ordered logit model as well price change predictions from an ordered logit with one additional discrete price change of \$0.50 and an ordered logit with additional discrete changes of \$0.50, -\$2 and +\$4. The addition of the \$0.50 price change does little to change the predicted price changes. Adding larger price changes of -\$2 and +\$4 leads to larger predicted price changes for larger tax increases, particularly for tax increases above roughly \$0.55. While these larger price changes would shift the burden of bigger tax increases towards consumers and increase the social cost of taxes, they do not change the general shape of the predicted price change curve.

Finally, we assess the sensitivity of the results to the exclusion of the change in wholesale price polynomial. Figure E6 plots the change in price and pass-through rate predicted by two ordered logits with a quartic polynomial in the tax rate. One is the baseline model presented in the text and includes all of the control variables, x_{jt}; the second excludes all terms of the change in wholesale price cubic polynomial. The coefficients on the Δw polynomial terms in the baseline model are large in magnitude and significant in the ordered logit model. Dropping the Δw terms reduces the fit of the model as the AIC changes by $5,942,783.71 - 5,912,234.04 = 30,549.67$.

However, when we exclude these terms and examine the key outputs of the ordered logit model in Figure E6 below, we find that there is little to no change in the relevant predictions. Once we average over all of the covariates x_{jt}, we find out that the average partial effects are relatively insensitive to the other included covariates. It is reassuring that the relationship between Δp and $\Delta \tau$ appears to be robust and not dependent on other covariates (at least once we average over x_{jt}). Even the apparent difference in implied pass-through rates is small except for at very small tax changes when a denominator approaching zero magnifies even very small differences.

F. Pass-Through by Tax Per Bottle

In Figure 6 we report summary pass-through measures averaged over all products and reported in terms of tax per liter. Below we break-out implied pass-through rates by bottle size. In general pass-through rates for different tax increases are similar across bottle sizes. This is as we would expect since model predictions are nearly identical for ΔP as a function of $\Delta \tau$ by size. The discrepancy in predicted price changes comes from differences in other state variables such as relative prices and wholesale prices, which don't look tremendously different on average across package sizes.

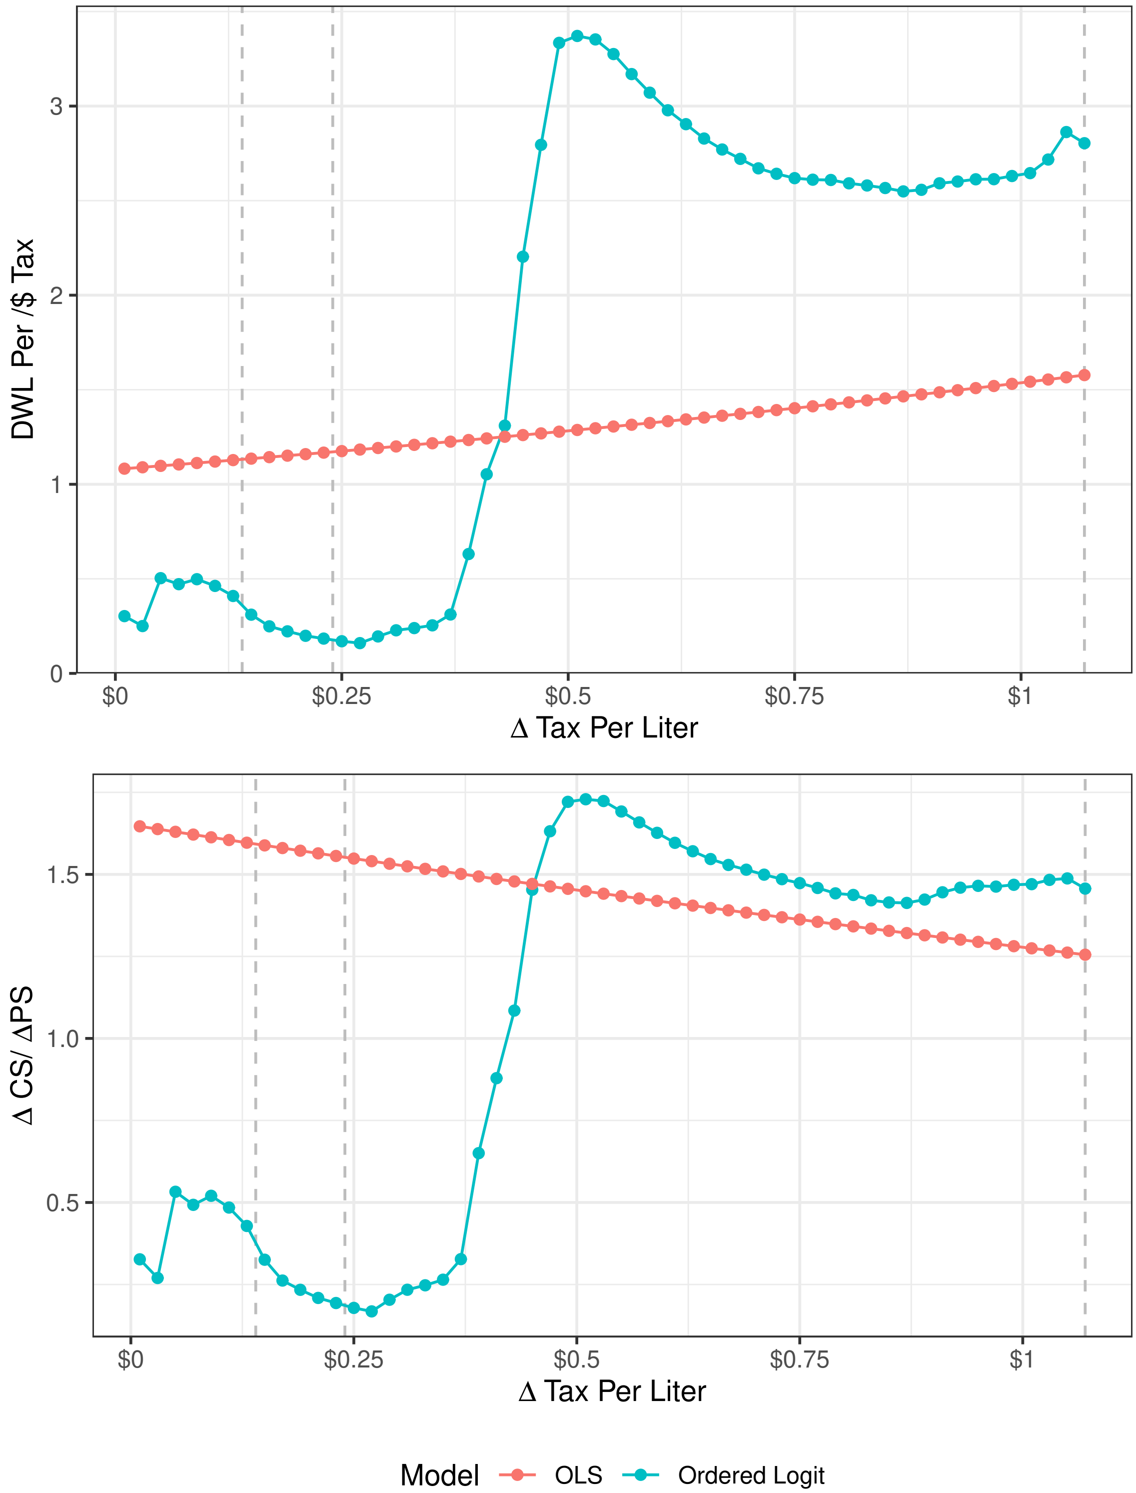

Figure E1: Welfare Predictions: Ordered Logit vs. OLS
Top Pane: Efficiency: DWL per Dollar of Tax Revenue; Bottom Pane: Incidence $\Delta CS/\Delta PS$.
Own Elasticity $\epsilon = -2.5$
Vertical Lines at Observed Tax Changes.

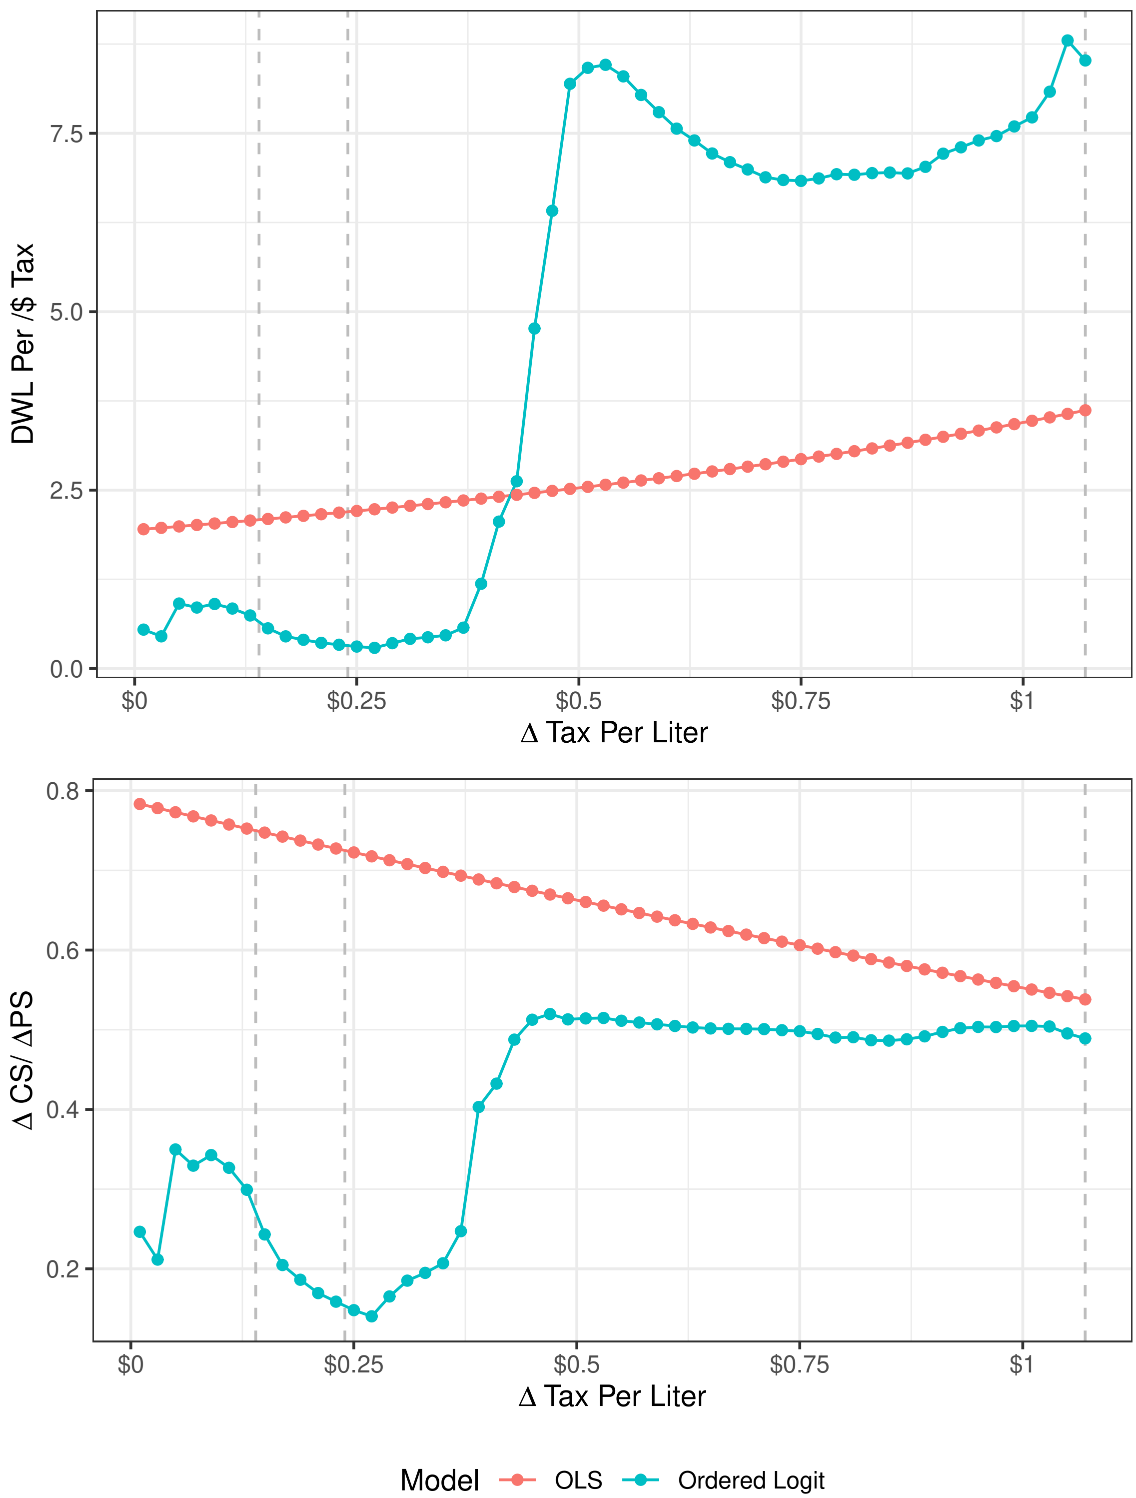

Figure E2: Welfare Predictions: Ordered Logit vs. OLS
Top Pane: Efficiency: DWL per Dollar of Tax Revenue; Bottom Pane: Incidence $\Delta CS/\Delta PS$.
Own Elasticity $\epsilon = -4.5$
Vertical Lines at Observed Tax Changes.

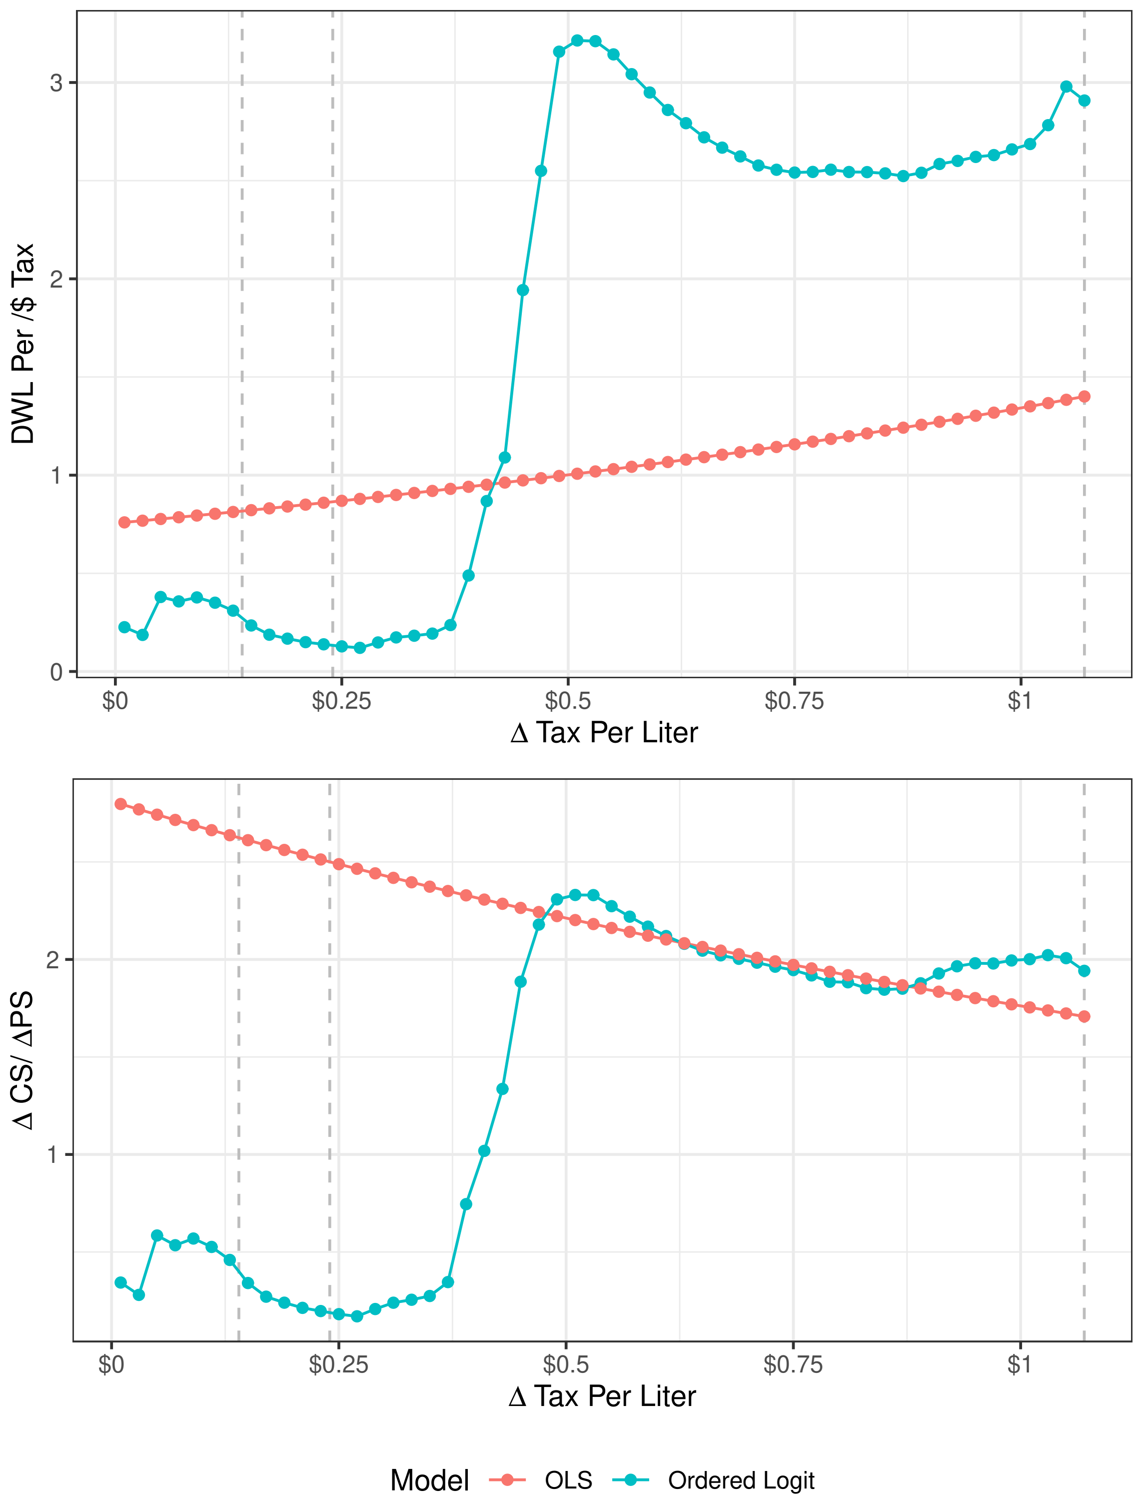

Figure E3: Alternative Markups $\mu = 1.2$: Ordered Logit vs. OLS
Top Pane: Efficiency: DWL per Dollar of Tax Revenue; Bottom Pane: Incidence $\Delta CS/\Delta PS$.
Own Elasticity $\epsilon = -3.5$
Vertical Lines at Observed Tax Changes.

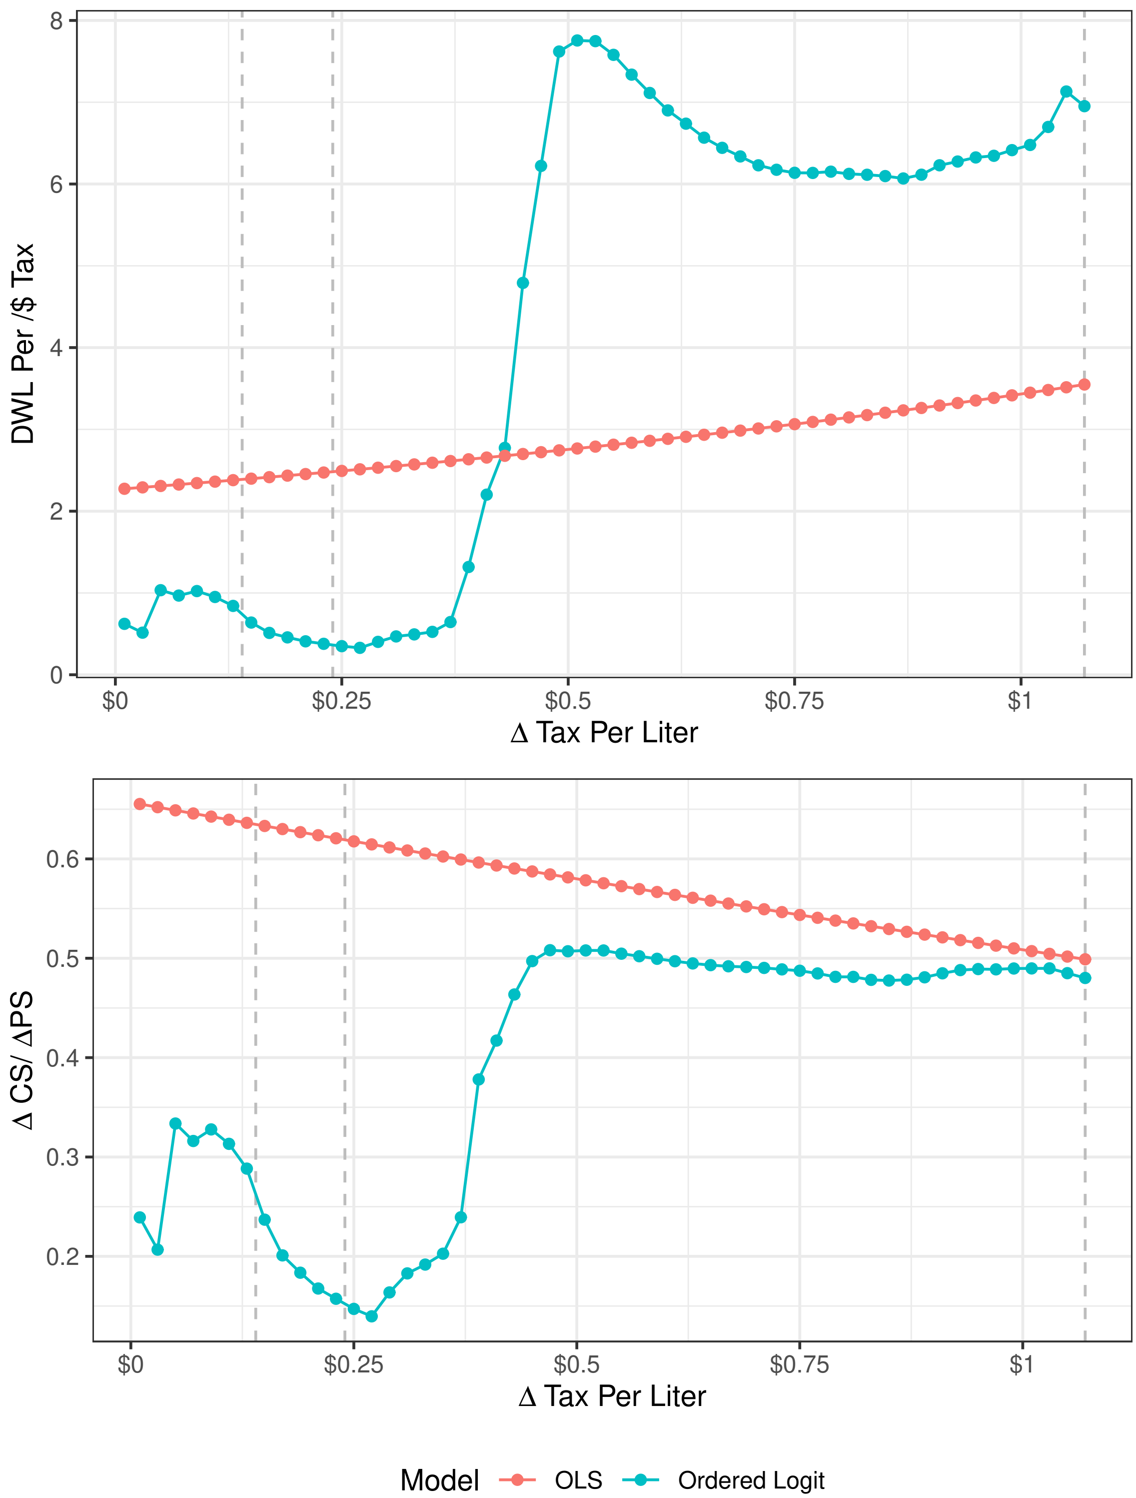

Figure E4: Alternative Markups $\mu = 2$: Ordered Logit vs. OLS
Top Pane: Effiecioncy: DWL per Dollar of Tax Revenue; Bottom Pane: Incidence $\Delta CS/\Delta PS$.
Own Elasticity $\epsilon = -3.5$
Vertical Lines at Observed Tax Changes.

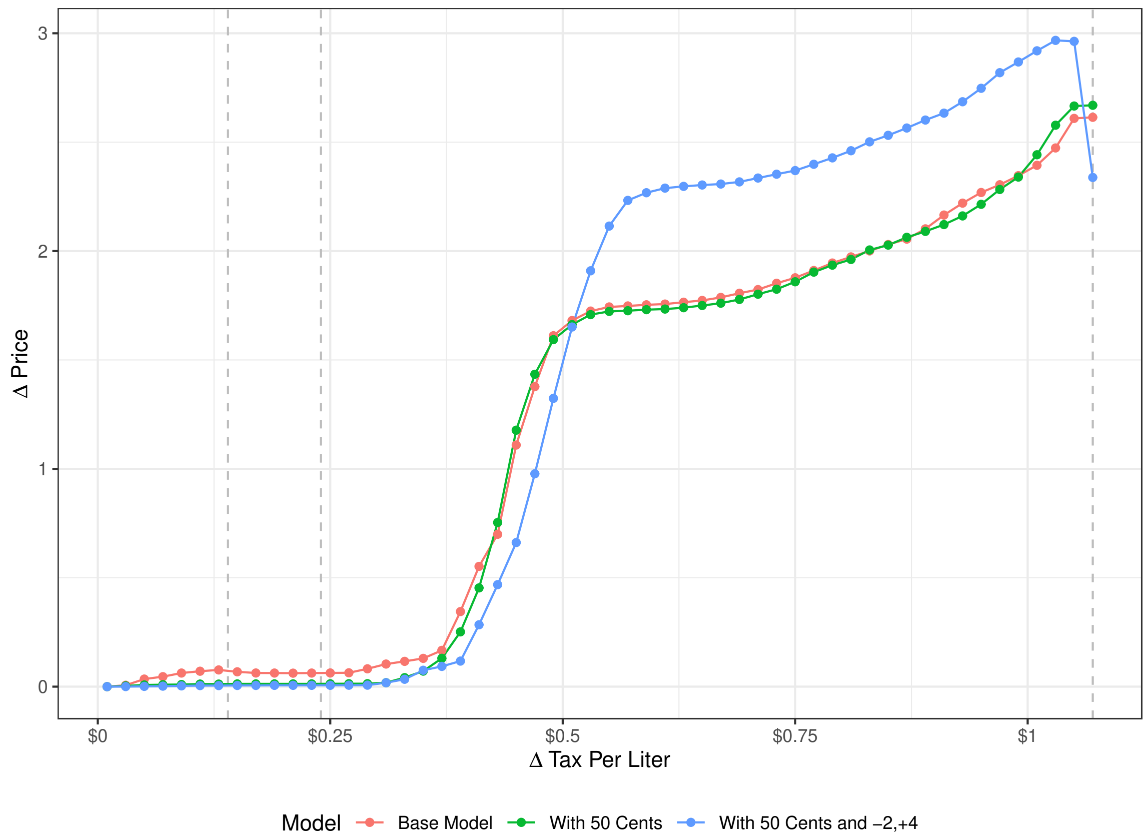

Figure E5: Robustness to Additional Price Points

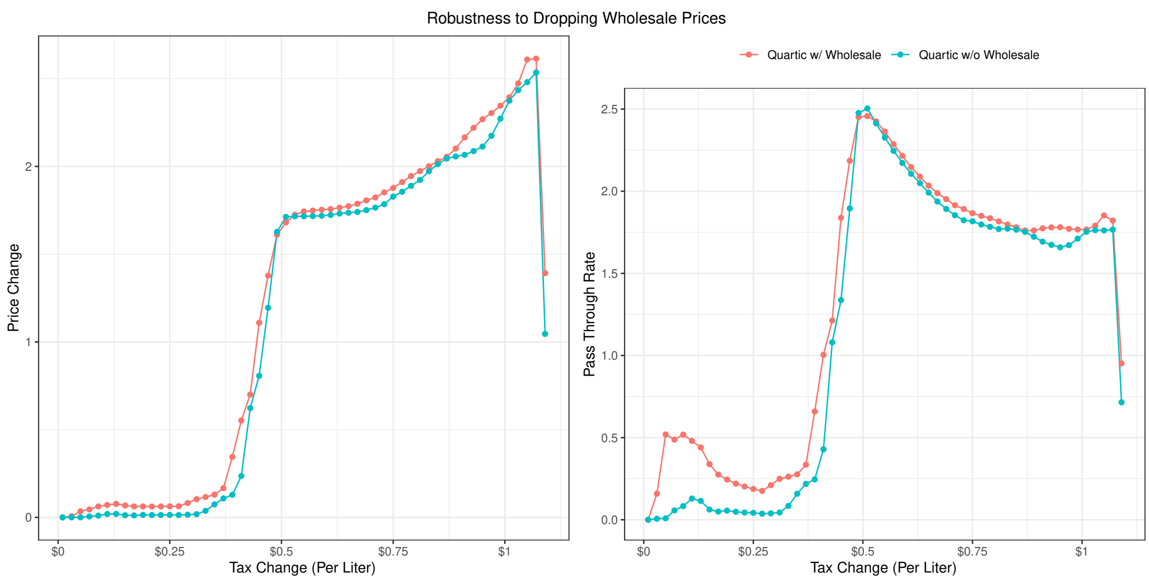

Figure E6: Robustness to Dropping Wholesale Price Terms

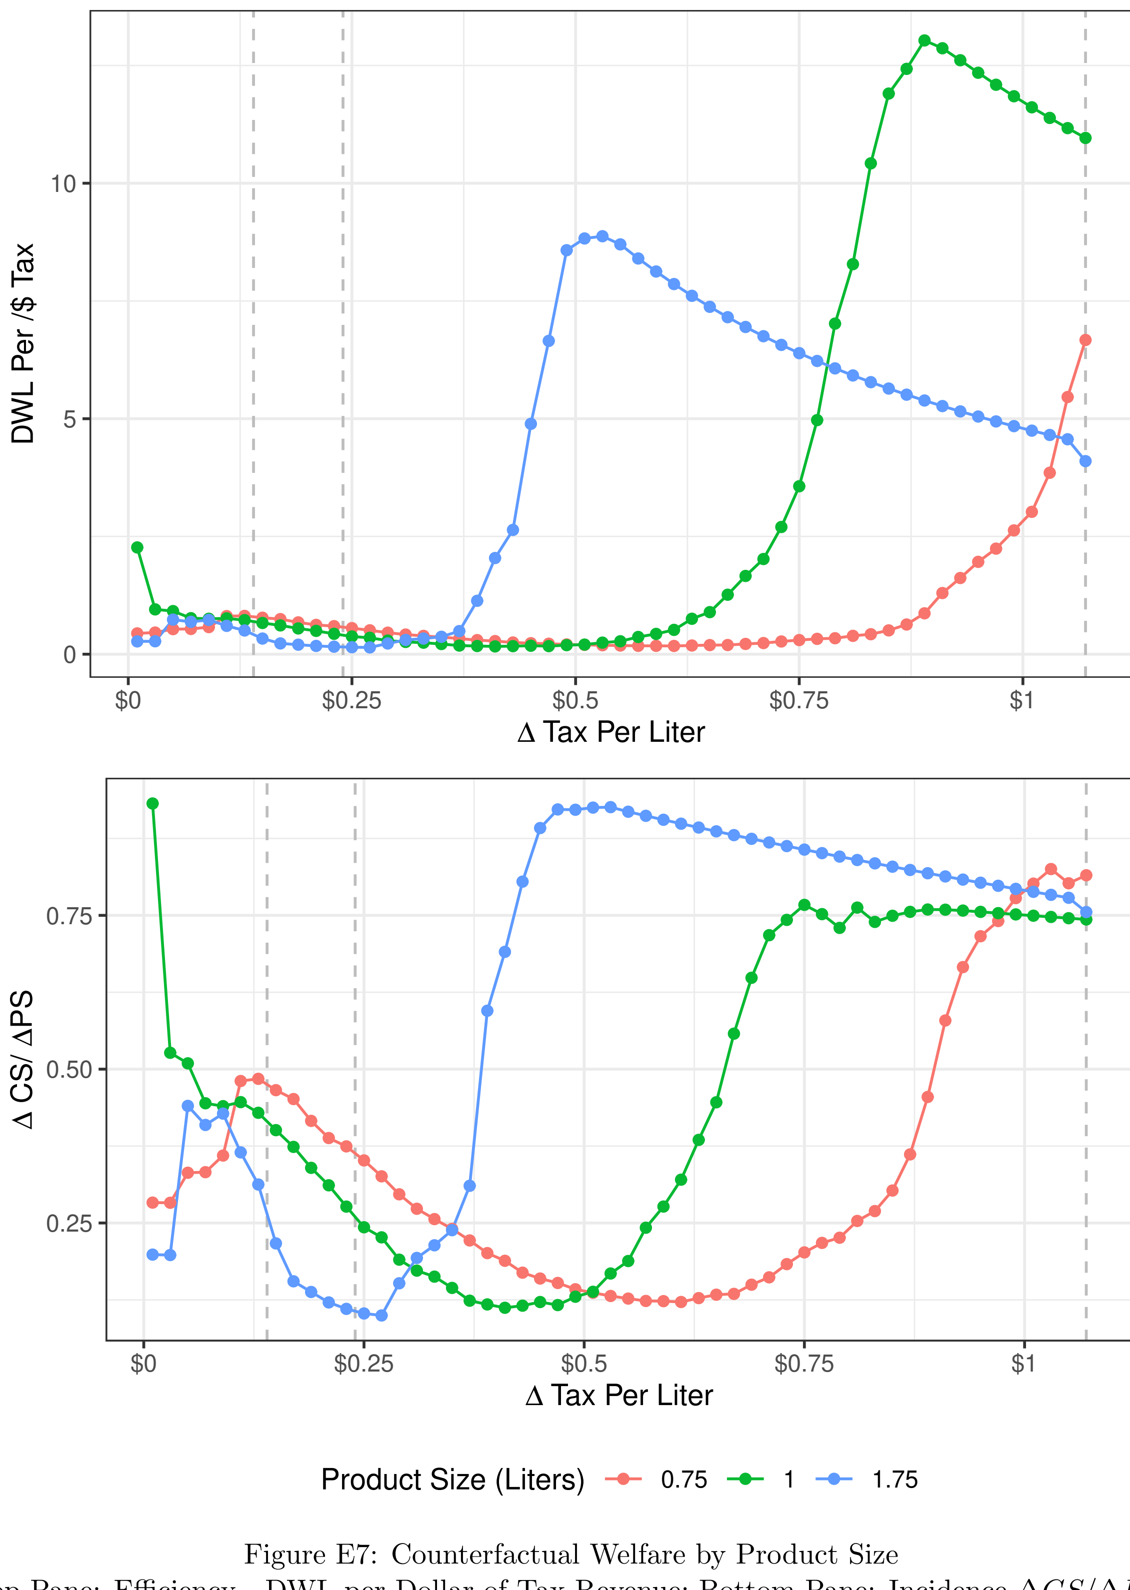

Figure E7: Counterfactual Welfare by Product Size
Top Pane: Efficiency - DWL per Dollar of Tax Revenue; Bottom Pane: Incidence $\Delta CS/\Delta PS$.
Own Elasticity $\epsilon = -3.5$
Vertical Lines at Observed Tax Changes.

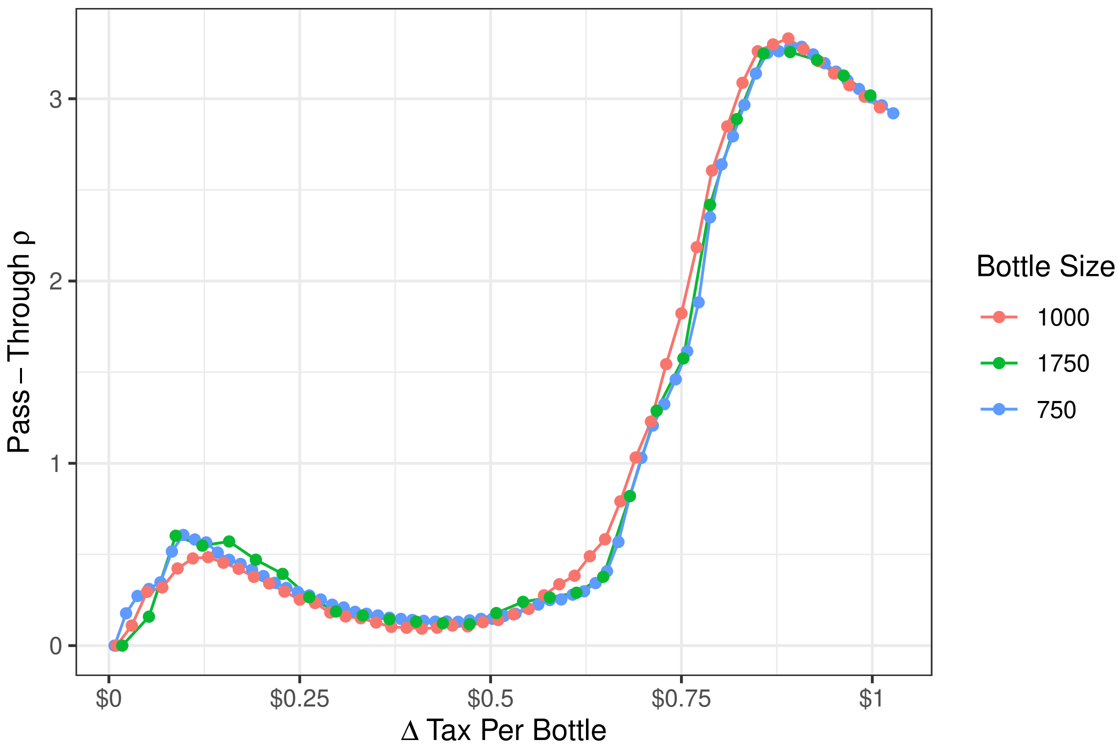

Figure F1: Counterfactual Welfare by Product Size
Implied Pass-Through Rate by Bottle Size
Vertical Lines at Observed Tax Changes.

G. Unabbreviated Tables

Table 5 reports only the key coefficients of interest from the estimated ordered logit models. The table below reports coefficients for all variables included in the regressions:

Table G1: Ordered Logit Estimates $\Delta p \in \{-1, 0, +1, +2, +3\}$

	Cubic	Quartic	Quintic	Spline(1)
Tax Change	493.006*	521.487*	521.816*	2.126*
	(28.187)	(28.033)	(28.003)	(0.539)
Tax Change^2	−99.167*	−129.960*	−128.773*	−8.467*
	(27.933)	(28.605)	(28.620)	(3.270)
Tax Change^3	−49.017*	−28.363	−36.295	36.666*
	(16.972)	(17.416)	(18.676)	(9.129)
Tax Change^4		−48.285*	−37.171	3.219*
		(13.532)	(17.072)	(0.235)
Tax Change^5			−12.768	
			(14.558)	
Wholesale Price Change = 0	−0.417*	−0.421*	−0.422*	−0.421*
	(0.152)	(0.153)	(0.154)	(0.153)
Wholesale Price Change	63.168	63.714	63.906	63.767
	(27.791)	(26.328)	(26.283)	(26.307)
Wholesale Price Change^2	−56.618	−37.501	−34.991	−36.644
	(22.336)	(21.878)	(22.283)	(21.964)
Wholesale Price Change^3	2.873	2.339	2.285	2.325
	(23.936)	(22.657)	(22.632)	(22.641)
Total Product Sales	−0.074	0.001	0.007	0.003
	(0.062)	(0.059)	(0.058)	(0.059)
Total Store Sales	0.196*	0.214*	0.215*	0.214*
	(0.040)	(0.038)	(0.038)	(0.038)
log Lag Price	−0.159	−0.113	−0.105	−0.110
	(0.100)	(0.110)	(0.112)	(0.110)
High Price	−0.201	−0.240	−0.248	−0.243
	(0.100)	(0.102)	(0.103)	(0.102)
Low Price	0.468*	0.485*	0.488*	0.486*
	(0.121)	(0.126)	(0.126)	(0.126)
Relative Price	−320.500*	−303.031*	−291.879*	−299.844*
	(7.079)	(7.295)	(7.286)	(7.346)
Relative Price^2	57.171	−2.198	−4.360	−3.377
	(25.652)	(26.608)	(26.394)	(26.455)
Relative Price^3	4.773	35.559*	35.594*	35.702*
	(10.220)	(10.208)	(10.322)	(10.323)
IL	1.885*	1.874*	1.889*	1.878*
	(0.616)	(0.636)	(0.641)	(0.638)
LA	0.233	0.205	0.322	0.235
	(0.431)	(0.449)	(0.471)	(0.450)
Total Prod Sales × IL	−0.160	−0.224	−0.229	−0.226
	(0.102)	(0.100)	(0.099)	(0.099)
Total Prod Sales × LA	0.093	0.039	0.020	0.033
	(0.085)	(0.082)	(0.082)	(0.082)
Total Store Sales × IL	−0.274*	−0.305*	−0.307*	−0.306*
	(0.084)	(0.081)	(0.081)	(0.081)
Total Store Sales × LA	−0.253*	−0.281*	−0.285*	−0.282*
	(0.058)	(0.056)	(0.056)	(0.056)
$\log(p_{j,t-1})$ × IL	−0.722*	−0.724*	−0.730*	−0.726*
	(0.212)	(0.218)	(0.219)	(0.218)
$\log(p_{j,t-1})$ × LA	−0.155	−0.167	−0.202	−0.176
	(0.150)	(0.156)	(0.161)	(0.156)
High Price × IL	0.196	0.225	0.233	0.227
	(0.170)	(0.172)	(0.172)	(0.171)
High Price × LA	−0.496*	−0.484*	−0.460*	−0.478*
	(0.145)	(0.147)	(0.150)	(0.147)
Low Price × IL	−0.446	−0.482	−0.486	−0.484
	(0.257)	(0.261)	(0.261)	(0.261)

Table G1 – continued from previous page	Cubic	Quartic	Quintic	Spline(1)
Low Price × LA	−0.106	−0.148	−0.134	−0.145
	(0.164)	(0.168)	(0.166)	(0.168)
Relative p × IL	−403.341*	−440.460*	−452.278*	−443.933*
	(4.654)	(4.818)	(4.768)	(4.816)
Relative p × LA	−56.814*	−82.062*	−87.300*	−83.808*
	(2.855)	(2.961)	(2.968)	(2.968)
Relative p^2 × IL	18.378	87.400*	89.939*	88.757*
	(14.275)	(14.790)	(14.608)	(14.657)
Relative p^2 × LA	71.773*	138.899*	136.407*	138.955*
	(8.702)	(9.072)	(9.028)	(9.043)
Relative p^3 × IL	16.511*	−20.770*	−20.896*	−20.983*
	(5.023)	(5.050)	(5.005)	(5.066)
Relative p^3 × LA	139.343*	117.293*	115.563*	116.913*
	(2.790)	(2.808)	(2.847)	(2.823)
Observations	2,371,792	2,371,792	2,371,792	2,371,792
State-UPCs	3,567	3,567	3,567	3,567
Out of Sample Likelihood	729,224	726,580	726,364	726,138
BIC	1,459,324	1,454,063	1,453,658	1,453,179

The table above reports estimates from ordered logistic regressions of quarterly price changes on quarterly tax changes with different parameterizations of the tax change and a number of controls. The first column employs a cubic orthogonal polynomial of the tax change while columns 2 and 3 use quartic and quintic orthogonal polynomials of the tax change, respectively. The final column uses a spline with a knot point at tax change = 1 and its coefficients are not polynomial order. The controls measure the change in wholesale price since the last change in retail price, annual sales of the product at that retailer, annual unit sales at that retailer of all products, the natural log of the price for the product the prior quarter at the same store, whether that store sold the product at the highest or lowest price the prior quarter and the difference between the price last quarter and the median competitor price for the same product last quarter. All four regressions are weighted by product-store sales in the year prior to the tax change. Weights are balanced by state, bottle size and tax change indicator. * Significant at the 1 percent level. All standard errors are clustered at state-UPC level.

References

Berry, S., J. Levinsohn, and A. Pakes (2004): "Automobile Prices in Market Equilibrium," *Journal of Political Economy*, 112, 68–105.

Conlon, C. T. and N. S. Rao (2019): "The Price of Liquor is Too Damn High: State Facilitated Collusion and the Implications for Taxes," Working Paper.

Miravete, E. J., K. Seim, and J. Thurk (2018): "Market Power and the Laffer Curve," *Econometrica*, 86, 1651–1687.