The Great Recession and Consumer Demand for Alcohol: A Dynamic Panel-Data Analysis of U.S. Households

September 25, 2014

Chad Cotti[a], Richard A. Dunn[b], and Nathan Tefft[c]

Abstract: For those looking to design policies that mitigate the deleterious consequences of alcohol abuse, understanding how consumer demand for alcohol responds to changes in the local economic conditions is of great importance. We use high-frequency purchase data from a large panel of U.S. households between 2004 and 2011 to examine how alcohol demand changes over the business cycle in a dynamic panel-data estimation framework. We find strong evidence that demand for packaged alcohol is procyclical. Changes in the state-level unemployment rate and personal per capita income between the most recent business cycle peak and trough imply a 6.5% decrease in the demand for packaged alcohol (ethanol by volume). The results also show that the decline in alcohol expenditures is primarily due to a decrease in quantity rather than an overall decrease in the price per ounce of ethanol purchased. Moreover, we improve on the related literature methodologically by accounting for consumption dynamics, as long-run demand for alcohol may differ from short-run demand because of habit formation in the quantity and type of alcohol consumed. Our results also indicate that failing to account for consumption dynamics will tend to understate the long-run association between macroeconomic conditions and alcohol demand.

JEL classification codes: I1, J68, E32

Keywords: alcohol demand, unemployment, macroeconomic conditions, business cycle

[a] Department of Economics, College of Business, University of Wisconsin Oshkosh, Oshkosh, WI 54901, USA

[b] Department of Agricultural and Resource Economics, College of Agriculture and Natural Resources, University of Connecticut, Storrs, CT 06268, USA

[c] Department of Economics, Bates College, Lewiston, ME 04240, USA

The authors thank Resul Cesur and seminar participants in the HPM Colloquium in Health Services Research at Yale University, the Population Health Sciences Seminar at the University of Wisconsin-Madison, and at the 2012 Southern Economic Association Annual Meeting for helpful comments. The authors thank the Kilts-Nielsen Data Center at The University of Chicago Booth School of Business for providing the data (http://research.chicagobooth.edu/nielsen/). Author order is alphabetic and lead authorship is shared amongst all of the authors. There are no conflicts of interest.

1. Introduction

The Great Recession has provided a once-in-a-generation opportunity to study how health behaviors and outcomes are affected by widespread, long-lasting macroeconomic shocks. At the height of the Great Recession in the United States, unemployment peaked at 10.0 percent (October 2009), the highest it had been since the 1982-83 recession, when unemployment rose to 10.8 percent (December 1982). Although peak unemployment during the Great Recession was slightly lower, economic recovery also progressed much more slowly. Between December 1982 and December 1983, the unemployment rate fell 2.5 points and dropped below 7.5 percent seventeen months after its peak. In contrast, between October 2009 and October 2010, the unemployment rate fell only 0.5 points and forty-five months passed before it fell below 7.5 percent.

The extended period of high unemployment associated with the Great Recession therefore presents a natural experiment that can be used to better identify both the short-run and long-run effects of macroeconomic conditions on health behaviors that exhibit habit formation. When utility depends on both current and past consumption, shocks that improve health behaviors today may result in better health behaviors in the future. Yet, the issues of consumption dynamics and long-run versus short-run responses have been largely ignored in the existing literature that investigates the relationship between macroeconomic performance and population health. For example, numerous studies have found that alcohol demand declines when the unemployment rate increases (Ruhm, 1995; Freeman, 1999; Dee, 2001; Ruhm and Black, 2002; Charles and DeCicca, 2008). But, none of these have accounted for consumption dynamics in their empirical specifications. As a result, they provide potentially inconsistent estimates of two distinct phenomena: 1) the association between alcohol demand and macro-economic conditions over the business cycle and 2) the long-run effect of a permanently higher or lower unemployment rate on alcohol demand.

Therefore, we use the Nielsen Homescan Consumer Panel Dataset (hereafter NHCPD) to estimate a model of dynamic demand consistent with habit formation spanning. Because the NHCPD is a continuous panel of American households with high-frequency purchasing data from the years before, during, and after the Great Recession, it permits examination of how alcohol demand responds to the macro-economic environment in both the short-run and long-run. Moreover, it allows us to examine two margins over which households could alter their alcohol purchasing habits, quantity (ounces of alcohol) and quality (price per ounce of alcohol). In contrast, previous studies have generally relied on alcohol sales or consumption data that are either cross-sectional at the individual level or aggregated to form state-level panels.[1]

In a preview of our results, we find strong evidence that demand for packaged alcohol is procyclical with a high degree of demand persistence. At both the aggregate and household levels, increases in the state-level unemployment rate are negatively related to quantity demanded of off-premises alcohol. The results from the household-level analysis are robust to the choice of covariates and we find procyclical demand dominates across all major alcohol subtypes. The results also show that the decline in alcohol expenditures is due to a decrease in both quantity and price rather than an overall decrease in the price per ounce of ethanol purchased. Finally, the negative relationship between the unemployment rate and alcohol demand is not explained by a decrease in total household grocery expenditures.

2. Data

To study how state-level measures of macroeconomic performance are associated with household alcohol purchases, we use data from the Nielsen Homescan Consumer Panel Dataset (NHCPD) between 2004 and 2011. The NHCPD sample includes respondents from all states and major metropolitan areas which generates a demographically and geographically balanced group of respondents. Nielsen assigns a projection factor to each household based on household characteristics including race, education level, income, age, presence of children, and household size so that nationally, regionally, and locally (market area) representative projections can be calculated. The dataset contains approximately 40,000 households from 2004 to 2007, and 60,000 households from 2007 onward.

Because the subsequent analysis estimates a flexible dynamic panel data model, it is important that the NHCPD includes households that report purchase information for several

[1] The work by Dávalos, Fang and French (2011), which uses the two waves of NESARC to estimate alcohol use, is a notable exception. Yet, a two-period panel with waves spaced 4 years apart would not be sufficient to examine the particular problem of dynamic demand.

years. The Nielsen Corporation provides respondents with incentives to encourage continued participation but not to influence purchasing habits. As a result, approximately 80% of panelists from one year continue to participate in the following year. For estimation purposes, this implies that the lagged first-difference can be calculated for roughly two-thirds of the sample and the twice lagged first-difference can be calculated for nearly half the sample.[2]

The Nielsen Corporation provides each participating household with a device to scan the UPC code of every item they purchase on every shopping trip. Each unique UPC code is treated as a separate item. For example, a six-pack of 12oz beer cans is coded as distinct from a six-pack of 12oz beer bottles. Thus, a single beer brand may have nearly a hundred UPC codes based on unit size, packaging, and number of units. If the consumer indicates making the purchase at a store that participates in Nielsen's point-of-sale data (POS) collection program, the item is assigned the average weekly price of that product at that store. If not, the NHCPD panelist is asked to provide the price. Participants are instructed to scan items from all shopping trips at all types of retail stores, including liquor stores, wine shops and beer distributors, not just shopping trips to supermarkets. The NHCPD does not, however, contain information on purchases at restaurants, bars or other on-premises consumption locations.

To construct our dataset, we first calculate total ounces of beer purchased by the household over all shopping trips in each month, as well as expenditure on beer products. For example, if for shopping trips during January 2007, we observe that a household reported purchasing a six-pack of 12oz *Budweiser* cans for $5 and a twelve-pack of 12oz *Miller Lite* bottles for $11, the household is assigned 12*6 + 12*12 = 216oz of beer for that month and $16 in beer expenditures. We then repeat this calculation for purchases of wine and for purchases of liquor. From these values we calculate the total alcohol by volume (ABV) purchased and total alcohol expenditure using the ethanol conversion factors employed by the Alcohol Epidemiologic Data System (AEDS): 0.045 for beer, 0.129 for wine, and 0.411 for spirits (LaVallee & Yi, 2012).

An immediate concern when using any data related to alcohol purchase and/or consumption is under-reporting, particularly in retrospective studies (Ramstedt 2010; Gmel and

[2] For purposes of comparison, 2.6 percent of Consumer Expenditure Survey (CEX) participants who complete their diary in week one do not return a diary in week two. The 80% retention rate in the NHCPD after 52 weeks implies a weekly attrition rate of 0.4 percent.

Daeppen 2007). Although the NHCPD is not retrospective, participants have the choice of scanning or not scanning any item they purchase. A comprehensive validation study has found that the reporting of quantities purchased in the NHCPD is 94% to 99% correct (Einav, Leibtag, and Nevo 2008), but this study did not look at disaggregated product groups, i.e. alcohol purchases, specifically.

To examine the degree of under-reporting in NHCPD relative to other household-level data sources, we compare expenditure on packaged alcohol for at-home use purchased by participants in the NHCPD and Consumer Expenditure Survey (CEX). The CEX utilizes a different data collection method than the NHCPD (expenditure diary rather than a UPC scanner) and CEX participation is for only two weeks. Panels A and B of Figure 1 plots per household weekly expenditure on beer and wine[3] among all respondents smoothed using a five-week moving average for the NHCPD sample period. The mean household weekly expenditure on beer is $1.18 for NHCPD participants compared to $2.45 among CEX participants, and mean expenditure on wine is $0.92 and $1.97, respectively.

Although participating households in the NCHPD tend to under-report relative to CEX households, a more careful consideration of the data reveals that this is largely because a greater proportion of households in the former never report purchasing alcohol. When restricting attention to households conditional on reporting positive purchases, however, the NHCPD and CEX are quite similar. Panels C and D of Figure 1 show that their respective time series track each other closely. The means of household weekly expenditures on beer, conditional on reporting alcohol purchases, are $8.82 for NHCPD participants and $9.09 for CEX respondents. For wine, the corresponding values are $6.71 and $7.06.

A related concern is that the nature of NHCPD data collection leads households to not report when the marginal value of time is high. Were this the case, under-reporting could vary systematically with macro-economic conditions. Although it is not possible to test this in an absolute sense, we can again examine whether the NHCPD is any worse in this regard than the

[3] We restricted attention to beer and wine categories, which are both clearly defined in CEX and NHCPD. In CEX, there is not a single "liquor" category recorded, but instead could be represented by several residual categories after beer and wine are excluded, in conjunction with a separate "whiskey" category. This lack of clarity could yield misclassification error with respect to expenditure comparisons.

CEX. From Figure 1, it is clear that while the NHCPD series follows the same general pattern of the CEX series, it is significantly less volatile. This possibly reflects the large sample of the former and suggests that employing the NHCPD will generate a higher degree of precision than using the CEX.

More formally, for every month between January 2004 and December 2011, we calculate per household expenditure of off-premises alcohol in both the NHCPD and CEX. We also calculate the proportion of households reporting any purchase of off-premises alcohol. We then estimate:

$$logA_t^{NHCPD} - logA_t^{CEX} = \beta_0 + U_t\beta + \varepsilon_{it}$$

where A_t^D is the alcohol purchasing measure (per household expenditure and proportion reporting any purchase) of interest in dataset *D*. If under-reporting was more likely to occur in NHCPD relative to the CEX during good economic times because it requires much more time to complete, then β should be positive. For both measures, however, we estimate that β is zero—not just statistically insignificant, but a well-estimated precise zero.

Based on these comparisons, in combination with findings from existing studies on the reliability of quantity information when a positive quantity is reported (Einav, Leibtag, and Nevo 2008), our baseline analysis excludes NCHPD households that never report any alcohol purchases, so that our sample suffers from under-reporting that is no worse than a comparable sample from the CEX. As robustness checks, we also conducted the primary analyses reported in sections 4 and 5 using a) all NHCPD participant data and b) only those household-year observations with positive reported purchases, with the natural log of purchases as the dependent variable. The results are similar to those obtained when dropping households who never report any alcohol expenditures, and all of the qualitative conclusions are the same.

2.1. Purchases versus consumption

Although expenditure data often form the basis from which alcohol demand is studied, the purchase of alcohol in-and-of-itself is not a population health problem. Rather, we care about alcohol purchases in so far as they are a predicate to consumption. A number of studies have identified that individual purchasing behavior is strongly associated with future consumption behavior (Amlung et al. 2012; Murphy et al. 2009; MacKillop and Murphy 2007), including total consumption and risk of binge drinking.

Nevertheless, the NHCPD has two limitations that should be acknowledged: 1) it is not possible to assign household purchases to individuals and 2) it is not possible to assign purchased alcohol to particular drinking events. A large purchase of alcohol could reflect that a household is hosting a social event (alcohol is consumed—perhaps irresponsibly—by non-household members) or that the household is stockpiling alcohol because of a sale (alcohol is consumed responsibly over a long time period) or that a household member intends on binging for several days (alcohol is consumed dangerously). Each of these scenarios has different public health implications. Of course, these issues are not unique to the NCHPD. Indeed, Ruhm's seminal paper utilizes state-level excise tax receipts to calculate per capita consumption (Ruhm 1995) and thus faces the same criticism. Moreover, for the purpose at hand, it is doubtful that these issues influence estimation results, at least qualitatively. If a household increases alcohol purchases because they wish to host more parties, it is also likely they consume more alcohol at parties hosted by other households. For many households, building an inventory of alcohol is an invitation to increase consumption, as well. Hence, though it is necessary to admit the shortcomings of expenditure data, we believe it is still highly informative about consumption, and therefore a valuable empirical tool.

3. Nielsen Scantrack market-level analysis

3.1. Methods

Although the primary focus of this paper is to analyze household purchasing behavior over the business cycle in the presence of consumption dynamics, we first consider aggregate market-level changes in alcohol purchases to specifically compare the Nielsen data with earlier studies in an attempt to isolate that the differences in our primary findings (presented in section 5) are not a function of the data or time frame. For example, Ruhm (1995) studies state-level aggregate household purchases, and in this section we demonstrate that a very similar analysis using Nielsen's market-level projections yields broadly similar results.

Nielsen has defined 54 mutually exclusive geographic areas as Nielsen Scantrack markets ("markets") and provides projection factors that can be used to generate sample statistics that are representative at the market level. We therefore construct representative households for each market by calculating the weighted arithmetic mean of monthly consumption and expenditure amounts. Using an analogous method, we similarly construct weighted market-level means of annual consumption and expenditure.

We begin by evaluating how changes in macroeconomic conditions are associated with changes in different measures of alcohol demand. We adopt the standard estimation approach in the literature of a fixed-effects specification across the 54 Nielsen Scantrack markets. The estimating equation is:

(1) $$Y_{it} = \beta_0 + X_{it}\beta + \gamma U_{it} + \alpha_i + \tau_t + \varepsilon_{it}$$

where subscript i denotes Scantrack market and subscript t denotes time period, and all standard errors are corrected to allow for non-independence of observations from the same market through clustering (Bertrand, Duflo, and Mullainathan 2004). The terms α_i and τ_t represent market and time fixed effects. Y_{it} is the natural logarithm of per household alcohol purchases in a market-year, or market-month-year. Also, when estimating models of this type, it is common to use a weighted least squares approach, because the precision of the alcohol consumption estimates varies with the underlying group size (Ruhm 1995). Estimates in the market-level alcohol consumption models are therefore weighted by the square root of the market population.

The variable U_{it} is a measure of its market's overall unemployment rate constructed as the population-weighted average of the state-by-month unemployment rates for each county contained in the Scantrack market. Thus, for this initial specification, γ is the approximate percentage increase in per household alcohol purchases over the defined duration (month or year) that accompanies a one percentage point increase in the market unemployment rate.[4] It is worth noting that in the presence of habit formation, which is likely a salient aspect of alcohol demand, there is no *a priori* reason to expect that the estimated coefficient on unemployment

[4] The exact percentage change in per household alcohol by volume purchased is exp(γ)-1.

in equation 1 using annual data will equal the effect calculated from the estimated coefficient using monthly data.[5]

Following earlier work, we recognize that this empirical strategy should isolate the impact of changes in the unemployment rate from other determinants of market-level alcohol consumption. We therefore include X_{it}, which is a vector of covariates that includes market per capita personal income, beer taxes, and area population, all constructed by similarly aggregating from each market's underlying counties. State-by-month unemployment rates are drawn from the Local Area Unemployment Statistics of the Bureau of Labor Statistics, U.S. Department of Labor[6], state-by-year beer tax rates were downloaded from the Tax Foundation[7], and state-by-year personal per capita income are from the U.S. Department of Commerce Bureau of Economic Analysis, Regional Economic Accounts.[8]

A notable omission from the set of explanatory variables is a measure of alcohol price. Ruhm (1995) argues strongly that because price is endogenous, it should not be included in the set of explanatory variables. If higher unemployment reduces demand, this will cause suppliers to reduce the price. To the extent that price changes that are correlated with the unemployment rate are demand driven (and they very likely are during the time period under consideration) we want the coefficients on macroeconomic conditions to capture this process.

Table 1 reports summary statistics for the relevant variables aggregated by market. Evident in Table 1 is that the estimation timeframe spans a period of substantial variation in state-by-state macroeconomic conditions. For example, unemployment rates fluctuate across market-years from a low of 2.33% to a high of 14.8%.

---

[5] As an example, suppose that the level alcohol demand follows a simple autoregressive process:

$$Y_{it}=40+.8Y_{it-1}-2u_{it},$$

so that the steady-state level of alcohol consumption at an unemployment level of 5% is 150. First, assume that beginning from steady-state the unemployment rate is constant at 5% for 12 months, then increases by 0.2 percentage points each month for 12 months, then decreases by 0.2 percentage points each month for 12 months. An OLS regression of log annual alcohol demand on mean unemployment rate yields a coefficient estimate of -0.057 compared to a coefficient estimate of -0.051 using log monthly demand. Assume instead that beginning from steady-state the unemployment rate is constant at 5% for 12 months, then increases by 0.2 percentage points per month for 24 months. Annual data yields an estimate of -0.059 compared to -0.061 using monthly data.

[6] http://www.bls.gov/lau/home.htm

[7] http://taxfoundation.org/

[8] http://www.bea.gov/regional/index.htm

3.2. Results

We begin by estimating equation (1) for a sample where market-level purchases have been aggregated annually and do not control for variation in area population as a direct comparison with Ruhm (1995). Results are reported in Panel 1 of Table 2. Column (1) provides results where the dependent variable used to measure alcohol consumption is the total alcohol-by-volume of ethanol (ABV) purchased in a market each year. Using this measure for total off-premises alcohol consumption yields results that are consistent with Ruhm (1995), showing that increases in the unemployment rate lead to declines in overall alcohol consumption. Columns (2) through (4) look at how consumption of specific types of alcohol is impacted by changes in macroeconomic conditions. Specifically, we find that, while all estimates are negative, only liquor sales are statistically impacted by changes in economic conditions in a meaningful way, again matching Ruhm (1995), who also found that after controlling for area per capita personal income, only liquor (spirits) was sensitive to changes in the area unemployment rate. While on the one hand it should not be unexpected to find similar results, it is somewhat surprising and reassuring to find results from two different studies, which use data that is 20-30 years apart, have different measures of alcohol demand, and are aggregated at different levels (state vs. market), yet demonstrate qualitatively identical findings.

Next we investigate how annual aggregation plays a role in the outcomes presented in Panel 1. Specifically, the Nielsen data allow for monthly aggregation of the alcohol measures. Estimates presented in Panel 2 are re-estimates of equation (1) where the monthly aggregation allows for the inclusion of year-month period dummies. Controls for variation in market population are also included. Results of this re-estimation are presented in Panel 2 of Table 2 and the findings are extremely consistent with those presented above and the general findings of previous area-level analysis of a negative response of alcohol consumption to increases in state unemployment rates.

4. Household-level analysis

4.1. Purchasing patterns of households

The primary contribution and focus of this paper is a study of how individual households' alcohol purchasing behaviors are associated with adverse economic conditions as proxied by the state unemployment rate and personal per capita income. Tables 3 and 4 present summaries of average annual alcohol purchases, average annual prices, and annually reported demographic characteristics for sample household-year observations between 2004 and 2009. Households that never purchased any alcohol (17.7%) were excluded, as discussed in section 2. The summarized household-years are only those observations for which lagged ABV may be calculated, in order to maintain the same sample throughout the subsequent regression estimation procedures.

Table 3 reports annual purchases by volume (oz), dollars ($), price per oz ($), and proportions of positive purchases for ABV and the three alcohol subcategories beer, wine, and liquor.[9] Of the estimation sample, approximately 23% of household-years involve no alcohol purchases. Average prices are assigned according to Nielsen's POS reporting program or reported by households, as described above, for households with at least one purchase in a given year. Since the reported prices also account for adjustments to list price including discounts, coupons, etc, some of the prices for a positive number of purchases are equal to zero (only 0.1% of prices in the estimation sample).[10]

Next, Table 4 includes a summary of demographic characteristics for household-years in the estimation sample. All of the reported demographic characteristics are recorded as categorical variables in the NHCPD, so we only report percent values. Since purchases are recorded at the household level, demographics are organized either by household or by male or female household head. The sample is evidently not representative of the U.S. population because groups such as household heads who are under age 25 or who do not have a high school degree are under-sampled. Because we are primarily interested in within-household changes in alcohol purchases, however, this limitation will not present an issue unless there is

[9] We also conducted the main analysis after dropping purchases that represented greater than 7 ounces of ABV per day (0.4% of the sample), or the equivalent of 12 bottles or cans of beer per day. These results were very similar to when using the full sample.

[10] Omitting observations with a zero price does not influence the main regression results.

substantial heterogeneity in responses to certain population subgroups conditional on the observed demographic characteristics.

Our final summary table for the household-level analysis (Table 5) shows area-level characteristics after merging into the Nielsen panel. Unemployment rates, per capita income, and beer tax rates are merged in as described in the market-level analysis section above. In the household-level analysis we also include quarterly, county-level supply-side controls for the number of establishments of several types that sell alcohol, as well as total employment in those establishment categories. These data are from the Quarterly Census of Employment and Wages of the Bureau of Labor Statistics, U.S. Department of Labor.[11]

4.2 Empirical specification

The need to account for consumption dynamics has long been recognized in demand analyses. Pollak (1970) argued that long-run demand may differ from short-run demand for three reasons: consumers may face short-run commitments that make changes in demand prohibitively expensive; consumers may have limited information regarding possible substitute goods that prevents them from adjusting demand immediately; and consumer preferences may exhibit habit formation, so that current preferences depend upon past consumption.

Habit formation is likely to be a particularly salient aspect of alcohol consumption and previous studies tend to find strong evidence for habit persistence (Baltagi and Griffin, 1995, 2002; Grossman, Chaloupka, and Sirtalan, 1998; Freeman, 2000; Eakins and Gallagher, 2003). For example, a negative life-event may cause an individual to greatly increase her alcohol consumption. As a result, she develops alcohol dependency and maintains a higher level of alcohol consumption even when the shock disappears. More generally, individuals may develop a taste for a particular level of alcohol consumption or a particular type of alcohol that persists over time.

Building upon the theoretical results of Pollack (1970), assume that households maximize the following Gorman form utility function:

$$U_h^t(X_h^t) = \sum_{k=1}^{K} a_k \, log(x_{hk}^t - b_{hk}^t) \qquad a_j > 1, \left(x_{hj}^t - b_{hj}^t\right) > 0, \sum_{k=1}^{K} a_k = 1,$$

[11] http://www.bls.gov/cew/

subject to the period budget constraint:

$$\sum_{k=1}^{K} p_k^t x_{hk}^t < m_h^t,$$

where *i* denotes household, *t* denotes time, and *k* denotes good. For each good:

$$b_{hj}^t = B_{hj}^t + \beta_j x_{hj}^{t-1},$$

so that B_{hj} denotes the pre-committed consumption of good *j* and β_j capture the psychological effect of habit formation. To simplify the analysis, further assume that $b_{hj}=0$ for all $j \neq alcohol$.[12] Demand for alcohol, *A*, would then be:

$$q_{hA}^t(q_{hA}^{t-1}, p^t, m_h^t) = B_{hA}^t + a_A \frac{m_h^t}{p_A^t} + a_A \beta_A x_{hA}^{t-1}.$$

The specification of B_{hA}^t therefore determines the demand relationship to be estimated, permitting a very flexible framework in which to introduce heterogeneity across household and time period. Specifically, let B_{hA}^t be a linear function of both observable, *X*, and unobservable, ρ, attributes:

$$B_{hA}^t = X_h^t \theta + \rho_i^t$$

The observable attributes that affect alcohol demand include household demographics, local area characteristics, and macroeconomic indicators. The unobserved attributes must be specified to capture variation across time and household that would otherwise bias the coefficient estimates of interest. Obvious candidates include household, month, and geographic fixed-effects along with an idiosyncratic random component. But, high frequency models of alcohol demand raises two important additional considerations.

First, shocks that cause individuals to change the level or the composition of alcohol consumption have persistent effects even in the absence of habit formation. For example, the death of a loved one may result in psychological harm that lasts for several periods, causing a prolonged increase in alcohol consumption. This suggest that the idiosyncratic error should follow an MA process. Second, high frequency panel data may exhibit household-specific cyclical patterns. For example, all households may increase their alcohol consumption during the winter holiday season, but the inclusion of month fixed-effects that are common to all

[12] This assumption ignores the real possibility that alcohol demand depends upon the demand of other habit-forming goods, such as gambling and smoking. Future studies that examined how demand for all three goods interacted over the business cycle could be a valuable extension to the results presented subsequently.

households will fail to adequately capture that some households systematically tend to increase their consumption more than other households. Instead, the effect of household-specific seasonality would incorrectly pass into the lagged dependent variable, tending to overstate the importance long-run dynamics and introducing an additional source of endogeneity. These aspects of alcohol consumption motivate the following empirical specification:

(2) $$A_{hst} = A_{hst-1}\beta_1 + M_{st}\beta_M + X_{hst}\beta_X + D_{st}\beta_D + \tau_{hm} + \tau_y + \gamma_s + \delta_h + \varepsilon_{hst} + \rho\varepsilon_{hst-1},$$

where A_{hst} is the amount of alcohol by volume (ABV) purchased by household *h* in geographic area *s* at time (year-month) t. ABV is measured in levels, rather than in logs, because households commonly report zero alcohol purchases in a given month. Hence, estimating in levels is necessary to maintain a comparable sample across different specifications when using household-level data.

M_{st} includes macroeconomic indicators, e.g. the unemployment rate and per capita income. A unique feature of the recovery from the Great Recession is the slowly declining unemployment rate that reflects both a decline in labor force participation among the long-term unemployed, as well an increase in employment. Thus, as a robustness check, we also add the market-area employment-to-population ratio to M_{st}.

X_{hst} contains household- and household head-level demographic characteristics (e.g. income, employment status, etc.). Again, because alcohol price is endogenous, it is not included in the set of explanatory variables. In addition, alcohol consumption could influence labor-market outcomes, introducing endogeneity. We lack good instruments for contemporaneous employment status and household, therefore, we use their lagged values to mitigate this problem and serve as general measures of socio-economic status.[13] D_{st} includes geographic area controls expected to affect individual alcohol consumption such as area-level tax rates and the availability of alcohol retailers, i.e., the number of liquor stores and bars, as well as employment in such establishments.

τ_{hm} is a household-specific month fixed-effect to address potential household-specific cyclicality; τ_y is a common year-fixed effect; γ_s is a vector of indicator variables for state of

[13] Given the dynamic demand framework, using lags of employment status and income will tend to reduce the potential effect of endogeneity, but not completely eliminate it. Nevertheless, we do not think the endogeneity of these variables will meaningfully influence the coefficient estimates of interest.

residence; δ_h are household fixed effects; and $\varepsilon_{hst} + \rho\varepsilon_{hst-1}$ is the composite MA(1) error term. Of course, higher order lags of both the dependent variable and the error term could be added to the right-hand side of equation (2), but as we report subsequently, specification tests reject the need for their inclusion.

It is worth noting that although panel-Tobit procedures exist (Honoré 1993), the orthogonality conditions required for method of moments estimation require that the unobservable in period *t* is independent of the unobservable in period *t-1*. The assumption of an MA(1) error in equation (2) violates this condition. As we report subsequently, the Arellano-Bond specification tests strongly reject that the composite error is uncorrelated over time in both the annual and monthly analysis, indicating that the MA(1) structure is necessary for consistent estimation. Thus, we do not model purchases as a latent-variable process, which is consistent with the treatment of expenditure in the infrequency of purchases model (IPM) (Keen 1986). In the IPM, zeros arise because an action that does occur (purchasing alcohol during 2008) is not observed occurring within a given time period (during January 2008). At least in cross-sectional data, Monte Carlo simulation results that show OLS provides less biased estimates of marginal effects than either the Tobit or hurdle models when the data follow an IPM process (Keen 1986; Stewart 2013).

We consider two approaches to estimating the parameters in equation (2) when including household-specific seasonality: aggregating to annual purchases and seasonal-differencing monthly observations. Each is treated in turn.

4.3 Annual purchases

Aggregating equation (2) over time to reflect annual purchases yields the following estimating equation:

$$(3) \qquad A_{hsy} = \beta_0 + A_{hsy-1}\beta_1 + M_{sy}\beta_M + X_{hsy}\beta_X + D_{sy}\beta_D + \tau_y + \gamma_s + \tilde{\delta}_h + \varepsilon_{hsy} + \rho\varepsilon_{hsy-1},$$

where *y* denotes the year. When summed, the household-specific month fixed-effects become part of a composite household fixed-effect: $\tilde{\delta}_h = \delta_h + \sum_m \tau_{hm}$.

Equation 3 defines a dynamic panel data model (habit formation) with an MA(1) error process (persistent shocks). Because the lagged dependent variable, A_{hsy-1}, is a function of the lagged unobservable, ε_{hsy-1}, Equation 3 suffers from two sources of endogeneity. First, A_{hsy-1} is correlated with the composite unobservable term, $\varepsilon_{hsy} + \rho\varepsilon_{hsy-1}$. Second, inclusion of the household fixed-effect demeans all the variables, including the unobservable terms. The mean unobservable within a household is a function of ε_{hsy-1} and hence the demeaned unobservable, $\varepsilon_{hsy} - \overline{\varepsilon_{hs}}$, is correlated with the lagged dependent variable.

The first-differenced counterpart to Equation 3 is:

(4) $$\Delta A_{hsy} = \Delta A_{hsy-1}\beta_1 + \Delta M_{sy}\beta_M + \Delta X_{hsy}\beta_X + \Delta D_{sy}\beta_D + \widetilde{\tau_y} + \Delta\varepsilon_{hsy} + \rho\Delta\varepsilon_{hsy-1}.$$

Though the household fixed-effect is removed, eliminating the second source of endogeneity described previously, it does not address the first source. Indeed, by first-differencing, the lagged first difference, $\Delta A_{hsy-1} = A_{hsy-1} - A_{hsy-2}$, is correlated with two terms in the composite error: $(\rho - 1)\varepsilon_{hsy-1} - \rho\varepsilon_{hsy-2}$.

Building on the results of Arellano and Bover (1995), Blundell and Bond (1998) suggest estimating Equations 3 and 4 as a system with appropriate lags of the explanatory variables as instruments. Specifically, lags of the levels are used as instruments in the first-difference equation (3) and lags of the differences are used as instruments in the level equation (2). Looking at the level equation, the equation for ΔA_{hsy-2} is the highest lagged difference that is not a function of ε_{hst-1}. Thus, A_{hsy-1} is instrumented using all possible first-differences of A_{hsy} lagged two periods and greater: $\{\Delta A_{hsy-2}, \Delta A_{hsy-3}, \dots\}$. Looking at the difference equation, A_{hst-3} is the highest lagged level that is not a function of $\Delta\varepsilon_{hsy-1}$. Thus, ΔA_{hsy-1} is instrumented using all possible levels of A_{hsy} lagged three periods and greater: $\{A_{hsy-3}, A_{hsy-4}, \dots\}$.

4.4 Seasonal differences

Aggregating to annual data obviates one of the principal benefits of the NHCPD: access to high frequency data. In addition, the unemployment rate is constructed based on a monthly survey (the annual unemployment rate is simply the average of monthly unemployment over the calendar year). Recognizing these benefits to analyzing how monthly expenditures vary with the monthly unemployment rate, an alternative approach to estimating equation (2) is to seasonally difference the data. That is, to remove the household-specific month effect from equation (2), subtract its 12th lag. By doing so, equation (2) becomes:

$$(5)\quad A_{hst} - A_{hst-12} = (A_{hst-1} - A_{hst-13})\beta_1 + (M_{st} - M_{st-12})\beta_M + (X_{hst} - X_{hst-12})\beta_X + (D_{st} - D_{st-12})\beta_D + \tilde{\tau}_y + (\varepsilon_{hst} - \varepsilon_{hst-12}) + \rho(\varepsilon_{hst-1} - \varepsilon_{hst-13}),$$

Intuitively, equation 5 is identified by asking whether alcohol demand in one period is different than demand in the same period last year when unemployment is different than it was in the same period last year. Equation 5 suffers from endogeneity because the seasonal difference of lagged dependent variable is correlated with the composite error. Thus, we instrument $A_{hst-1} - A_{hst-13}$ with the seasonal difference of the second lag: $A_{hst-2} - A_{hst-14}$. Hence, to test if alcohol demand this November, which is higher than the preceding November, predicts whether alcohol demand this December will be higher than the preceding December, we generate exogenous variation by considering whether alcohol demand this October was higher than alcohol demand the preceding October.

5. Results

5.1 Annual quantity

Table 6 reports coefficient estimates from multiple regression models, culminating in estimates of equation 3 in the rightmost columns. Displaying results from these several models serves to a) demonstrate the influence of introducing additional model structure and covariates on the coefficients of interest and b) offer a comparison with the application of models used in the existing literature to the NHCPD. Column 1 reports estimates from an ordinary least squares (OLS) regression without covariates and with heteroscedasticity-robust standard errors, as a baseline. Here, higher state-level per capita income is significantly associated with a lower quantity of alcohol purchases, but the estimate on the state unemployment rate is insignificant and small, suggesting has no meaningful impact. Associations from this specification can only be interpreted as causal and informative if all unobserved characteristics are assumed to be uncorrelated with the error term, an untenable assumption.

Columns 2 and 3 apply methods similar to those used in the existing literature to address potential omitted variable bias, adding state and year fixed effects, as well as state trends, in column 2 (e.g. Ruhm, 2000) and separately adding household and year fixed effects in column 3 (e.g. Dávalos, Fang, and French, 2012). Standard errors are clustered at the state and household level, respectively. In both cases, when controlling for unobserved fixed characteristics the unemployment rate coefficient estimates are now negative and suggest a strong procyclical response to increases in the unemployment rate. The per capita income coefficients do not vary substantially, although estimates are now statistically insignificant in both cases.

To our knowledge, existing studies do not address the dynamics of household alcohol consumption, as discussed above, considering dynamics is both important for accurate inference and an issue which previous research have been unable to investigate because of limitations in data. We therefore report results from estimating the system defined by equations 3 and 4 in the remaining columns.[14] After controlling for lagged purchases, column 4 reveals that the dynamic panel estimate of the unemployment rate coefficient is substantially smaller of the estimate when ignoring consumption dynamics (compared to column 3). Thus, failing to account for dynamic demand behavior overstates the contemporaneous effect of a change in the unemployment rate.

But, there is also strong persistence in alcohol purchases, i.e. there is a positive and significant correlation between lagged and current ABV purchases. The long-run effect of a permanent, one percentage point increase in the unemployment rate is equal to $\beta_M/(1-\beta_1)$. Based on the coefficient estimates reported in column 4, the long-run effect of is -2.76. Thus, failing to account for consumption dynamics understates the long-run effect of a permanent change in unemployment.

Recognizing that the relationship between alcohol purchases and macro-shocks captured by the unemployment rate may not be linear, column 5 reports coefficient estimates when the unemployment rate-squared is added to the set of explanatory variables. Results of this specification show that both the level term and the squared term of the unemployment rate are significant, suggesting that the increases in the unemployment rate reduce that alcohol

[14] Using an Arellano-Bond test for zero correlation in the first-differenced errors, we confirmed that the lag structure of instruments identified for equations 3 and 4 above is appropriate. Specifically, zero correlation was rejected in the first lag (p-value 0.0000) but was not rejected in the second lag (p-value 0.9669).

purchases, but the slope of the relationship depends on the value of the unemployment rate. Specifically, these results suggest that a one percentage point increase in the unemployment rate from the mean value of approximately 7% is associated with a decrease of contemporaneous ABV of roughly 0.8 ounces, or approximately 0.65% of the sample mean ABV. The long-run effect of a permanent one percentage point increase from 7% to 8% is 4.4 ounces (3.6% of the sample mean). However, when evaluated at the national unemployment rate at the beginning of the recession of 4.8%, a one percentage point increase in the unemployment rate is associated with a decrease in contemporaneous ABV purchases of 1.6 ounces (8.8 ounces in the long-run), while at the peak unemployment rate of the last recession of approximately 10% there is no meaningful effect. These outcomes indicate that increases in the unemployment rate are associated with significant changes in alcohol purchases when the economy is in a relatively good state, but these effects begin to fade as the state of the economy worsens. Columns 6 and 7 report results when adding in time-varying covariates summarized in tables 4 and 5, and the results are very stable (Appendix Table A1 reports all of the coefficient estimates).[15]

In order to place the macroeconomic effects in the context of real events rather than a hypothetical change, we next calculated the total effect of changes in the unemployment rate and per capita income observed between the dates of the most recent business cycle's peak and trough (the peak-to-trough effect), reported as a separate row in table 6.[16] Between December 2007 and June 2009 the U.S. unemployment rate increased from 4.8% to 9.7%[17], and per capita income decreased from \$36,071 and \$35,496[18] (deflated to 2009 dollars). According to our results, the calculated peak-to-trough effect for the most recent business cycle is -7.8 ounces (p-value = 0.0000) in our preferred specification (column 7), which is very similar to the calculated long-run effect when evaluated at a base unemployment rate of 4.8% (8.73 ounces).

[15] Adding the employment-to-population ratio to the set of explanatory variables made the coefficient estimate on the unemployment rate slightly more negative, suggesting that while the decline in labor force participation during the recovery from the Great Recession attenuated the procyclicality of alcohol demand, it did not do so by a meaningful amount.

[16] The dates are assigned by the National Bureau of Economic Research (NBER): http://www.nber.org/cycles.html

[17] Data are reported from the U.S. Bureau of Labor Statistics, Labor Force Statistics from the Current Population Survey: http://www.bls.gov/cps/

[18] This is reported as national per capita disposable personal income by the Bureau of Economic Analysis National Income and Product Accounts Tables: http://www.bea.gov/national/index.htm

5.2 Monthly quantity

Table 7 offers a comparison of the two approaches outlined above in section 4. Column 1 replicates the results from column 7 of table 6 (our preferred annual dynamic panel-data model specification) for ease of comparison. Column 2 reports the results from estimating the monthly dynamic panel model with seasonally (same-month) differenced data in equation 5, instrumenting for the first lagged seasonal difference using the second lagged seasonal difference. Multiplying estimated effects by 12 yields the annualized seasonal difference impact: a one percentage point increase in the unemployment rate (evaluated at the mean) is associated with a 3.8 ounce decrease in annualized ABV purchases and a $1,000 increase in state per capita income is associated with a 2.4 ounce decrease in annual ABV purchases.

The contemporaneous annualized estimate from the dynamic panel-data model with monthly data (column 2) is more than twice that estimated in the dynamic panel-data model aggregated annually, and we find similar magnitudes for the long-run effect, which is equal to a 15.2 ounce decrease. The annualized peak-to-trough effect for the most recent business cycle is also larger, specifically a 17.94 ounce decrease in ABV purchased.

In general, the model using monthly data yields larger annualized estimates than does the dynamic panel-data model using data aggregated to annual quantities purchases. Although it is useful to confirm the estimated unemployment rate relationship using higher frequency data, in subsequent analyses we report results from the annual dynamic panel-data models only. This is primarily because the annual estimates are less likely to suffer from overestimating non-linear relationships when extrapolating from linear models (i.e. generalizing from estimates at a monthly level to the year or business cycle levels).

5.3 Alcohol subcategories

We use equation 2 to estimate demand for three types of alcohol—beer, wine, and liquor—and Table 8 shows estimation results from this analysis of alcohol subcategories.[19] The top panel

[19] Although it is straight-forward to derive a system of demand equations using the approach developed by Pollack (1970) when multiple goods exhibit habit formation, it is too restrictive for estimating the demand of alcohol subcategories. Specifically, it restricts past consumption of one good to be negatively related to current

reports estimates of effects on ABV in the first column (again duplicated from column 7 of table 6 for ease of comparison) and purchased ounces of beer, wine, and liquor in the last three columns. For both beer and wine there is a negative and statistically significant unemployment rate effect, while the state per capita income effect is statistically significant for beer and liquor (negative coefficient for the former and positive for the latter). For beer and wine, the peak-to-trough effect is negative and significant (p-value: 0.01). Evaluated at the mean unemployment rate, the largest per ounce marginal unemployment rate effect was for beer (-12.0), followed by wine (-7.3) and liquor (-0.4), but relative to the mean the largest effect is for wine (2.4% decline), followed by beer (1.1%) and liquor (0.2%). Note that liquor had the strongest association with the business cycle in the results from the market-level analysis, but now shows no significant response, which reinforces the importance of studying household-level behavior.

5.4 Expenditure versus quantity

In addition to quantity, it is also informative to study the effects on alcohol expenditures and inferred prices. The remainder of Table 8 reports coefficient estimates when the analysis is repeated replacing ABV with price per ounce of alcohol (in the second panel) and expenditures (in the third panel).

Although ABV price per ounce overall is negatively associated with unemployment rate changes, there are heterogeneous (offsetting) non-linear effects across categories. In particular, the price per ounce of beer and liquor generally increases with the unemployment rate, yielding a positive peak-to-trough effect (p-value: <0.0105). Although these effects are interesting, they are somewhat difficult to interpret because they do not offer information on whether households also adjust the quality of their alcohol purchases in response to changing macroeconomic conditions.

Lastly, the third panel of Table 8 reports the net changes in expenditures on each alcohol subcategory arising from changes in quantities and prices. From a mean annual

---

consumption of all other goods. As a robustness check, however, we also estimated equation 2 with demand for alcohol type *k* as a function of lagged demand of all types (not reported). The lagged cross-effects were not statistically significant and did not change the coefficient estimates on either lagged own-effect or macroeconomic indicators.

expenditure on alcohol from retail outlets of $169.28, a one percentage point increase in the unemployment rate (evaluated at the mean) yields a 3.0% decrease in contemporaneous expenditure. The long-run effect of a permanent one percentage point increase unemployment is a 19.8% decrease in expenditure on alcohol from retail outlets. For beer, wine, and liquor (mean annual expenditures of $61.68, $53.49, and $54.11) a one percentage point increase in the unemployment rate (evaluated at the mean) causes contemporaneous expenditures on alcohol demand from retail outlets to fall by 0.6%, 1.6%, and 0.8%, respectively.

5.5 Changes in quality

The regression results when price per ounce of ABV is the dependent variable suggest that households do not respond to changes in macroeconomic conditions by changing the quality of the alcohol purchased. To investigate this further, we calculated the mean price for every unique UPC associated with a 750ml package of gin, rum, table wine, tequila, vodka, and whiskey that was purchased during the 2004-2006 period.[20] Within each alcohol type, brands were ranked by their mean price. Brands above the median were categorized as *high quality*, while brands below the median were categorized as *low quality*. This categorization was then applied to all matching UPCs for all years in the NHCPD dataset. For example, the mean purchase price for a 750ml bottle of Bombay Sapphire Gin between 2004 and 2006 was $21.40, while the median price of all gins sold in a 750ml package was $9.23. Thus, Bombay Sapphire was categorized as high quality for all years and associated UPCs in the NHCPD dataset.

For each market area and each alcohol type, we calculated per household purchases of high quality and low quality brands (given the level of disaggregation by both type and quality, a household level analysis would exhibit an overwhelming majority of zero observations). We then re-estimated equation (1) for each alcohol-quality combination aggregated annually. These results (available upon request) reveal whether changes in unemployment shift purchases from higher to lower quality types. Overall, they show no evidence of a differential response by quality to changes in unemployment. For example, the two pooled models that

[20] 750ml is a standard volume of alcohol (see 27 CFR Chapter I, Part 5, Subpart E, Section 5.47a Metric standards of fill for distilled spirits bottled after December 31, 1979) and the most common package size for retail wine and spirits. It based on the older measure of a *fifth* (1/5th of a gallon or 757ml).

include purchases in all of the alcohol type categories (with a fixed effect added for each category) yield unemployment rates that are not statistically different from zero or each other. Of the twelve regression models run for each disaggregated alcohol-quality pair, only one unemployment rate coefficient was statistically significant at the ten percent level, hence providing very limited evidence of meaningful quality changes.

5.6 Comparison with all reported household expenditures

The NHCPD captures a large number of household purchases, so it is possible to compare all recorded expenditures with expenditures on alcohol as an assessment of whether alcohol purchases decline proportionally with total household expenditure during periods of adverse macroeconomic conditions. Table 9 replicates alcohol expenditure effects in the first column (drawn from the third panel of Table 8), and in the second column results are reported when replacing alcohol expenditures with a tally of all household expenditures recorded in NHCPD. As noted earlier, a one percentage point increase in the unemployment rate (evaluated at the mean) decreases expenditure by approximately $3.00, which represents a 1.8% expenditure decrease, while the $56.40 decrease (when evaluated at the mean unemployment rate) in all recorded expenditures represents a 0.7% decrease (the corresponding peak-to-trough effects are 29.81% and 0.55%). Again, these results should be interpreted with caution because not all household expenditures are recorded in the NHCPD, but they may be taken as suggestive evidence that alcohol is generally far more elastic than aggregate spending on groceries during economic downturns.

6. Discussion and Conclusion

In this paper, we used the Nielsen Homescan Consumer Panel Dataset to examine how off-premises alcohol purchases vary with macroeconomic indicators at the household-level. We find strong evidence that demand for packaged alcohol is procyclical. Specifically, we find that decreases in state-level unemployment are positively related to quantity demanded of packaged alcohol at the aggregate and household level. Our primary results from the household-level analysis are robust to the choice of estimation procedure, including dynamic panel design, and covariates selection. Results are also consistent across all major alcohol subtypes.

Moreover, our analysis adds meaningfully to the literature methodologically by accounting for consumption dynamics, as long-run demand for alcohol may differ from short-run demand because of habit formation in the quantity and type of alcohol consumed. Our results indicate that failing to account for consumption dynamics will tend to over-state the short-run association between macro-economic conditions and off-premises alcohol demand while understating the long-run association. This suggests that previous studies investigating the influence of recessions on alcohol consumption are reporting short-run effects that are smaller in magnitude than the long-run effects.

Further, regardless of estimation approach, our findings stand in contrast to the only other existing literature that uses within-person variation in alcohol consumption to study the relationship between recessions and alcohol abuse, by Dávalos, Fang, and French (2012). They find a robust positive relationship between binge drinking and the unemployment rate, which is not supported in the results reported here. Their analysis uses self-reports on binge drinking from wave 1 (2000-2001) and wave 2 (2004-2005) of NESARC. Thus, they only have 2 observations for each individual, whereas NHCPD allows us to construct a dataset with many more observations for each household, measured far more frequently. This will tend to increase precision and reduce the potential role of time-varying unobserved attributes.[21]

To put our findings in context, a one percentage point increase in the unemployment rate results in a long-run decrease of 8.7 ounces of alcohol per year purchased for off-premises consumption per alcohol consuming household, when evaluated at the national unemployment rate at the beginning of the great recession of 4.8%. This estimate is approximately 16.2 twelve-ounce cans of Budweiser beer. This would surely be considered a sizeable decline in consumer demand from a public health perspective.

[21] The NHCPD also includes more recent observations that span the collapse of the US housing market and subsequent financial crisis. In contrast, waves 1 and 2 of NESARC only encompass the relatively mild 2001 recession, and thus may not reflect how alcohol consumption responds to larger or more adverse fluctuations in macroeconomic conditions conditional on individual or household economic circumstances.

There are a number of limitations to the current analysis that warrant discussion. Firstly, any dataset that relies on self-reported alcohol consumption or expenditures will suffer from under-reporting, and the NHCPD is no different. Our investigation of the data suggests this largely arises from households choosing not to report any of their alcohol purchases. As a result, we chose to restrict the estimation sample to only those household reporting positive alcohol purchases at some point during each survey year. We thus have great confidence that our results are generalizable to the subpopulation of American adults willing to participate in the NHCPD and report their alcohol purchasing behavior. Although this group probably accounts for a large share of the US population—the Nielsen Corporation attempts to generate an NHCPD sample that is nationally representative and only 18% of households were omitted because they never reported purchasing alcohol—it must be acknowledged that the regression sample is not representative of the US population: it tends to undercount households with heads less than 25 years of age and households with low income. While it is reassuring that including all NHCPD participants in the analysis does not influence our estimates, extending our findings to younger and lower income households requires great caution. The key point to recognize, however, is that so many important questions about alcohol usage depend on self-reported data, that the alternative to accepting limits on generalizability is often no answer at all. With that in mind, we believe the preceding results provide valuable new information about demand for off-premises alcohol for a large segment of the US adult population.

A second important limitation of the NHCPD is its lack of information about purchases of alcohol at bars and restaurants. Though we are unable to test our hypothesis, given that on-premises alcohol is more expensive per unit of volume and our finding that households reduce total alcohol expenditures, we suspect that demand for on-premises alcohol consumption is more responsive to local macroeconomic conditions than off-premises alcohol. This is an issue that future research should address.

To summarize, this paper is the first to utilize high frequency data and a dynamic panel-data estimation framework to examine how alcohol demand changes over the business cycle. Results clearly demonstrate that off-premises alcohol consumption has a non-linear procyclical relationship and is generally consistent across sub-categories of alcohol (once consumption dynamics are accounted for appropriately). That said, a more complete understanding of how on- and off-premises alcohol purchases interact maybe important in determining an effective public policy response. Changes in off-premises purchasing behavior may be more salient for one subset of health shocks, e.g. domestic physical and sexual violence or falls in the home, while off-premises purchasing behavior may be more salient of other types of health shocks, e.g. the likelihood of involvement in a motor vehicle accident or being the victim of assault. Understanding these differences is of paramount importance, as the optimal public policy responses to these two different forms of consumption may vary greatly. We therefore consider this an important area of future research.

References

Amlung, Michael T, John Acker, Monika K Stojek, James G Murphy, and James MacKillop. 2012. "Is Talk 'Cheap'? An Initial Investigation of the Equivalence of Alcohol Purchase Task Performance for Hypothetical and Actual Rewards." *Alcoholism: Clinical and Experimental Research* 36 (4). Blackwell Publishing Ltd: 716–24. doi:10.1111/j.1530-0277.2011.01656.x.

Arellano, Manuel, and Olympia Bover. 1995. "Another Look at the Instrumental Variable Estimation of Error-Components Models." *Journal of Econometrics* 68 (1): 29–51. doi:http://dx.doi.org/10.1016/0304-4076(94)01642-D.

Baltagi, Badi H, and James M Griffin. 1995. "A Dynamic Demand Model for Liquor: The Case for Pooling." *The Review of Economics and Statistics* 77 (3). The MIT Press: 545–54. doi:10.2307/2109914.

———. 2002. "Rational Addiction to Alcohol: Panel Data Analysis of Liquor Consumption." *Health Economics* 11 (6). John Wiley & Sons, Ltd. 485–91. doi:10.1002/hec.748.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-In-Differences Estimates? ." *The Quarterly Journal of Economics* 119 (1): 249–75. doi:10.1162/003355304772839588.

Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics* 87 (1): 115–43. doi:http://dx.doi.org/10.1016/S0304-4076(98)00009-8.

Dávalos, María E, Hai Fang, and Michael T French. 2011. "EASING THE PAIN OF AN ECONOMIC DOWNTURN: MACROECONOMIC CONDITIONS AND EXCESSIVE ALCOHOL CONSUMPTION." *Health Economics*. John Wiley & Sons, Ltd. http://dx.doi.org/10.1002/hec.1788.

———. 2012. "EASING THE PAIN OF AN ECONOMIC DOWNTURN: MACROECONOMIC CONDITIONS AND EXCESSIVE ALCOHOL CONSUMPTION." *Health Economics* 21 (11). John Wiley & Sons, Ltd: 1318–35. http://dx.doi.org/10.1002/hec.1788.

Dee, Thomas S. 2001. "Alcohol Abuse and Economic Conditions: Evidence from Repeated Cross-Sections of Individual-Level Data." *Health Economics* 10 (3). John Wiley & Sons, Ltd. 257–70. http://dx.doi.org/10.1002/hec.588.

Eakins, John M., and Liam A. Gallagher. 2003. "Dynamic Almost Ideal Demand Systems: An Empirical Analysis of Alcohol Expenditure in Ireland." *Applied Economics* 35 (9). Routledge: 1025–36. doi:10.1080/0003684032000066796.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo. 2008. “Not-so-Classical Measurement Errors: A Validation Study of Homescan.” *National Bureau of Economic Research Working Paper Series* No. 14436. http://www.nber.org/papers/w14436.

Freeman, Donald G. 1999. “A Note on `Economic Conditions and Alcohol Problems’.” *Journal of Health Economics* 18 (5): 661–70. http://www.sciencedirect.com/science/article/pii/S0167629699000053.

———. 2000. “Alternative Panel Estimates of Alcohol Demand, Taxation, and the Business Cycle.” *Southern Economic Journal* 67 (2). Southern Economic Association: 325–44. doi:10.2307/1061473.

Gmel, Gerhard, and Jean-Bernard Daeppen. 2007. “Recall Bias for Seven-Day Recall Measurement of Alcohol Consumption Among Emergency Department Patients: Implications for Case-Crossover Designs.” *J. Stud. Alcohol Drugs* 68: 303–10.

Grossman, Michael, Frank J Chaloupka, and Ismail Sirtalan. 1998. “AN EMPIRICAL ANALYSIS OF ALCOHOL ADDICTION: RESULTS FROM THE MONITORING THE FUTURE PANELS.” *Economic Inquiry* 36 (1). Blackwell Publishing Ltd: 39–48. doi:10.1111/j.1465-7295.1998.tb01694.x.

Honoré, Bo E. 1993. “Orthogonality Conditions for Tobit Models with Fixed Effects and Lagged Dependent Variables.” *Journal of Econometrics* 59 (1-2): 35–61. doi:10.1016/0304-4076(93)90038-7.

Keen, Michael. 1986. “Zero Expenditures and the Estimation of Engel Curves.” *Journal of Applied Econometrics* 1 (3). Wiley Subscription Services, Inc., A Wiley Company: 277–86. doi:10.1002/jae.3950010305.

MacKillop, James, and James G Murphy. 2007. “A Behavioral Economic Measure of Demand for Alcohol Predicts Brief Intervention Outcomes.” *Drug and Alcohol Dependence* 89 (2-3): 227–33. doi:10.1016/j.drugalcdep.2007.01.002.

Murphy, James G, James MacKillop, Jessica R Skidmore, and Ashley A Pederson. 2009. “Reliability and Validity of a Demand Curve Measure of Alcohol Reinforcement.” *Experimental and Clinical Psychopharmacology*, 2009-23091-004, 17 (6). US: American Psychological Association: 396–404. doi:10.1037/a0017684.

Pollak, Robert A. 1970. “Habit Formation and Dynamic Demand Functions.” *Journal of Political Economy* 78 (4). The University of Chicago Press: 745–763 CR – Copyright © 1970 The University. doi:10.2307/1829929.

Ramstedt, Mats. 2010. “How Much Alcohol Do You Buy? A Comparison of Self-Reported Alcohol Purchases with Actual Sales.” *Addiction* 105 (4). Blackwell Publishing Ltd: 649–54. http://dx.doi.org/10.1111/j.1360-0443.2009.02839.x.

Ruhm, Christopher J. 1995. “Economic Conditions and Alcohol Problems.” *Journal of Health Economics* 14 (5): 583–603. http://www.sciencedirect.com/science/article/pii/0167629695000240.

———. 2000. “Are Recessions Good for Your Health?” *The Quarterly Journal of Economics* 115 (2): 617–50. doi:10.1162/003355300554872.

Ruhm, Christopher J, and William E Black. 2002. “Does Drinking Really Decrease in Bad Times?” *Journal of Health Economics* 21 (4): 659–78. http://www.sciencedirect.com/science/article/pii/S0167629602000334.

Stewart, Jay. 2013. “Tobit or Not Tobit?” *Journal of Economic and Social Measurement* 38 (3): 263–90. doi:10.3233/JEM-130376.

Figure 1. Per household weekly alcohol expenditure, five-week moving average

A. Beer expenditures

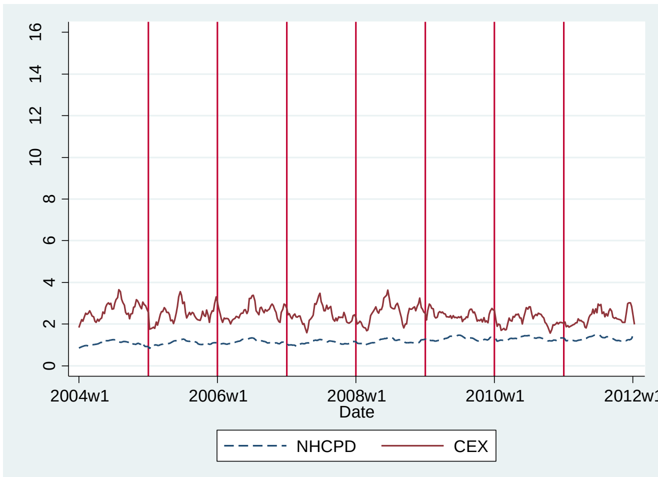

B. Wine expenditures

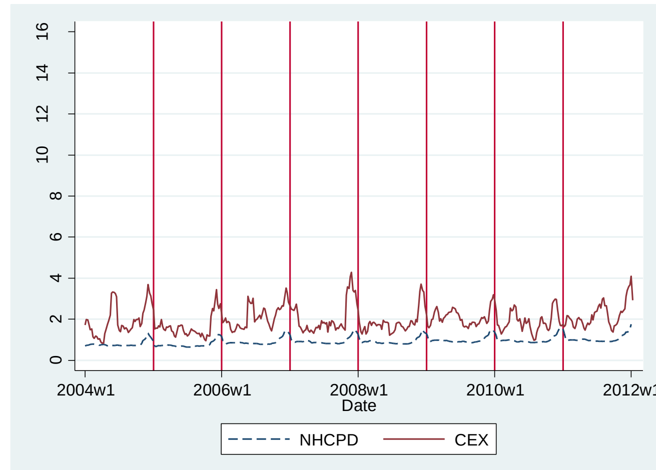

C. Beer expenditures, if alcohol expenditures > 0

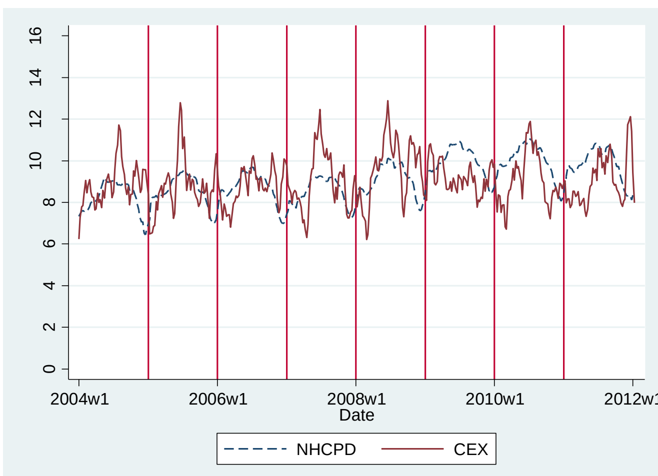

D. Wine expenditures, if alcohol expenditures > 0

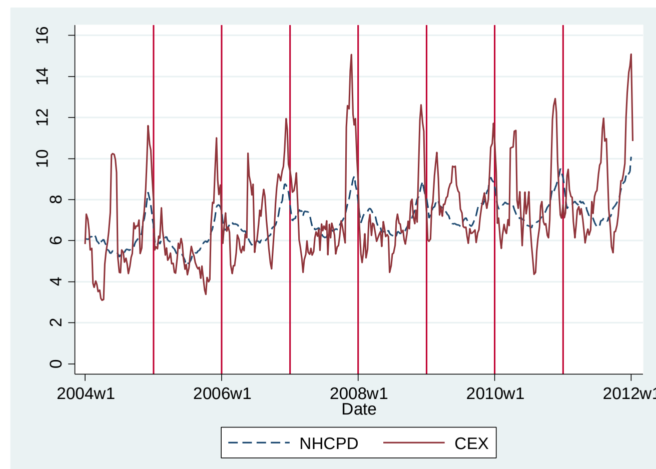

Table 1. Summary statistics, market-level analysis (*N* = 3,807)

	Mean	Std. Dev.	Min	Max
ABV per household (oz)	10.80	3.62	2.58	27.52
Beer per household (oz)	77.55	28.26	15.79	279.50
Wine per household (oz)	20.56	9.71	1.28	77.50
Liquor per household (oz)	11.34	5.14	1.20	41.01
Unemployment rate	6.81	2.51	2.33	14.80
Per capita income (2009 $s)	36750.04	8508.77	15870.02	57755.09
Beer tax per gallon	0.24	0.19	0.06	1.05
Population	4870421	2749782	1056565	17238362

Notes: Each measure is aggregated for 54 scantrack markets between 2004-2011 (not all markets are reportable for all years). ABV stands for alcohol by volume. Variables reported at the state level are included as county-level population weighted averages of the state-level values for all counties within a market area (markets often span states). All prices are deflated to 2009 dollars.

Table 2: Macroeconomic conditions and purchases per household, market-level analysis

	(1)	(2)	(3)	(4)
	Total ABV	Beer	Wine	Liquor
Panel 1. (N = 311)				
Unemployment Rate	-0.022**	-0.003	-0.015	-0.035**
	(0.009)	(0.013)	(0.013)	(0.014)
Panel 2. (N = 3,807)				
Unemployment Rate	-0.019***	-0.004	-0.015	-0.031***
	(0.005)	(0.009)	(0.010)	(0.008)

Notes: Each column in each panel represents a separate regression. Observations are aggregated by Nielsen scantrack market and year in Panel 1 and market and month in Panel 2. The dependent variable for all regressions is the natural log of the given alcohol measure. All regressions include controls for area deflated per capita personal income, beer taxes, location and time fixed effects, and were weighted by the square root of the market population. Regressions in Panel 2 also include controls for market area population. The standard errors in parentheses are corrected to allow for non-independence of observations within a market area through clustering. The sample includes observations from 2004–2011. ***,**,* denote statistical significance at the 0.01, 0.05 and 0.10 levels, respectively.

Table 3. Summary statistics, annual purchases, household-level analysis

	N	Mean	Std. Dev.	Min	Max
Beer (oz)	202,362	1049.35	3949.63	0	167534.00
Wine (oz)	202,362	304.00	1159.04	0	47001.46
Liquor (oz)	202,362	164.50	599.96	0	24599.62
ABV (oz)	202,362	154.04	371.51	0	10132.82
Beer ($)	202,362	61.68	206.77	0	7556.68
Wine ($)	202,362	53.49	181.63	0	7511.10
Liquor ($)	202,362	54.11	190.70	0	5991.07
Alcohol ($)	202,362	169.28	380.03	0	7734.17
Beer (oz) > 0	202,362	0.47	0.50	0	1
Wine (oz) > 0	202,362	0.53	0.50	0	1
Liquor (oz) > 0	202,362	0.41	0.49	0	1
ABV (oz) > 0	202,362	0.76	0.43	0	1
Beer price per oz. ($)	94,876	0.08	0.08	0	8.87
Wine price per oz. ($)	113,675	0.22	0.17	0	7.70
Liquor price per oz. ($)	82,378	0.42	0.34	0	20.80
ABV price per oz. ($)	154,976	1.37	1.42	0	197.05
All expenditures ($1000s)	202,362	8.36	5.42	0.44	89.79

Notes: Observations are aggregated by household and year. The estimation sample excludes households with no alcohol purchases (16.0%). ABV stands for alcohol by volume. Estimates are calculated using NHCPD sampling weights. Final price paid is calculated by subtracting reported coupons and discounts from the total price and then dividing by quantity, as detailed in the NHCPD documentation. Zero prices are not excluded because they are considered accurate net-of-discount prices by Nielsen. All prices are deflated to 2009 dollars.

Table 4. Summary statistics, head and household demographics

Binary variable	Mean
Female head present	78.3%
Male head present	70.7%
Female 25 ≤ Age < 55	46.5%
Female 55 ≤ Age < 65	15.1%
Female Age > 65	16.4%
Male 25 ≤ Age < 55	41.7%
Male 55 ≤ Age < 65	14.3%
Male Age > 65	14.5%
Female < college grad	52.8%
Female college grad	21.7%
Male < college grad	42.7%
Male college grad	23.1%
30 ≤ Female employ hrs < 35	3.8%
Female employ hrs ≥ 35	32.6%
Female not employed	33.6%
30 ≤ Male employ hrs < 35	1.7%
Male employ hrs ≥ 35	43.1%
Male not employed	23.5%
Household race black	11.9%
Household race oriental	3.3%
Household race other	8.8%
Household Hispanic	87.6%
Household widowed	11.4%
Household divorced/separated	19.8%
Household single	22.7%
$30k ≤ Household income < $60k	30.4%
Household income ≥ $60k	42.8%

Notes: *N* = 202,362. Observations are aggregated by household and year. The estimation sample excludes households with no alcohol purchases (16.0%). Estimates are calculated using NHCPD sampling weights. Omitted categories (not reported) are female employed < 30 hours, male employed < 30 hours, female age < 25, male age < 25, female < H.S. grad, male < H.S. grad, household white race, household not Hispanic, household married, household income < $30k. Reported male and female percents are unconditional means at the household level, i.e. these percents and omitted categories sum to 78.3% and 70.7%, respectively. See main text for head and household definitions.

Table 5. Summary statistics, state-by-year characteristics for the sample

	Mean	Std. Dev.	Min	Max
Unemployment rate	7.04	2.59	2.70	14.90
Per capita income ($1000s)	40.26	5.44	29.05	70.43
Beer tax per gallon	0.23	0.18	0.02	1.05
Beer, wine, & liquor stores	118.38	192.58	0	959.75
→ Employment	564.86	821.93	0	3511.67
Drinking places	162.72	215.47	0	1104.75
→ Employment	1437.04	1954.14	0	8406.25
Supermarkets	272.14	430.26	0.50	1886.50
→ Employment	9942.49	14651.18	0	72189.83
Convenience stores	108.98	138.25	0	615.00
→ Employment	577.76	814.88	0	4009.42

Notes: *N* = 202,362. Observations are aggregated by household and year (measures are reported at the state level). The estimation sample excludes households with no alcohol purchases (16.0%). Estimates are calculated using NHCPD sampling weights. All prices are deflated to 2009 dollars.

Table 6. State-level unemployment and annual alcohol purchases (ABV)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Unemployment rate	0.4050	-3.1137**	-2.3772***	-0.4821**	-3.3474***	-3.4684***	-3.0785***
	(0.3449)	(1.5240)	(0.8874)	(0.2289)	(1.0935)	(1.0984)	(1.0807)
Unemployment rate squared					0.1822***	0.1802***	0.1553**
					(0.0696)	(0.0691)	(0.0685)
Per capita income ($1000s)	-0.8887***	-1.1350	-0.5534	-0.0933	-0.0788	-0.1413*	0.0391
	(0.1476)	(1.1455)	(0.6349)	(0.0748)	(0.0750)	(0.0788)	(0.0854)
Lagged ABV (*t* - 1)				0.8257***	0.8218***	0.8212***	0.8222***
				(0.0329)	(0.0336)	(0.0336)	(0.0333)
Peak-to-trough effect	2.4955	-14.6045	-11.3301	-2.3086	-7.7863	-8.4373	-7.8018
Peak-to-trough p-value	0.1359	0.0641	0.0076	0.0316	0.0006	0.0002	0.0006
Model	OLS	State & Year FE & Trends	Household & Year FE	Dynamic PD	Dynamic PD	Dynamic PD	Dynamic PD
Standard Errors	Robust	Clustered (State)	Clustered (Household)	Clustered (Household)	Clustered (Household)	Clustered (Household)	Clustered (Household)
Household Demographics, Income, Employment	No	No	No	No	No	Yes	Yes
Area controls	No	No	No	No	No	No	Yes

Notes: N = 202,362. Observations are aggregated by household and year. For all models, the dependent variable is ABV (alcohol by volume) purchased. The estimation sample excludes households with no alcohol purchases (16.0%). A household fixed effects model that excluded households who moved between states (2.7%) and that implements standard errors clustered at the state level yields results nearly identical to those reported here (standard errors clustered at the household level).
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 7. Annual vs. Monthly Dynamic Panel Models (ABV)

	(1)	(2)
Unemployment rate	-3.0785***	-0.2803***
	(1.0807)	(0.0927)
Unemployment rate squared	0.1553**	-0.0026
	(0.0685)	(0.0058)
Per capita income ($1000s)	0.0391	-0.0204***
	(0.0854)	(0.0060)
Lagged ABV (*t* - 1)	0.8222***	0.7479***
	(0.0333)	(0.0551)
N	202,362	2,305,384
Peak-to-trough effect†	-7.8018	-17.9376
Peak-to-trough p-value	0.0006	0.0000
Model	Annual Dynamic PD	Monthly Dynamic PD (same-month differences)
Aggregation	Household-Year	Household-Month
Standard Errors	Clustered (Household)	Clustered (Household)
Household Demographics, Income, Employment	Yes	Yes
Area controls	Yes	Yes

Notes: Observations are aggregated by household and year in Column 1 and household and month in Column 2. For all models, the dependent variable is ABV (alcohol by volume) purchased. The estimation sample excludes households with no alcohol purchases. Unemployment rate and per capita income estimates for models with household-month aggregation show monthly change effects; multiplying by 12 yields an annualized effect comparable to the household-year estimates.

†The peak-to-trough effect in column 2 is multiplied by 12 so that both columns represent the annualized effect.

*** p<0.01, ** p<0.05, * p<0.1

Table 8. State-level unemployment and monthly alcohol purchases (subcategories)

	ABV	Beer	Wine	Liquor
Panel 1. Quantity (oz)				
Unemployment rate	-9.3439***	-56.1461***	-34.1070***	-6.3129
	(2.5001)	(20.7174)	(9.4060)	(4.5807)
Unemployment rate squared	0.5530***	3.1536**	1.9157***	0.4255
	(0.1605)	(1.3109)	(0.6181)	(0.2941)
Per capita income ($1000s)	0.0584	-2.5814**	-0.1088	0.4425***
	(0.0868)	(1.0426)	(0.3882)	(0.1604)
Lagged Dependent Variable (*t* - 1)	0.8176***	0.8150***	0.8375***	0.7898***
	(0.0341)	(0.0406)	(0.0519)	(0.0598)
Peak-to-trough effect	-19.81	-125.29	-76.95	-11.17
Peak-to-trough p-value	0.0000	0.0020	0.0000	0.2042
N	202,362	202,362	202,362	202,362
Panel 2. Price per oz ($)				
Unemployment rate	-0.0471**	0.0103***	0.0059	0.0330***
	(0.0202)	(0.0019)	(0.0060)	(0.0126)
Unemployment rate squared	0.0025*	-0.0006***	-0.0004	-0.0021***
	(0.0013)	(0.0001)	(0.0004)	(0.0008)
Per capita income ($1000s)	0.0012	0.0002**	0.0002	0.0021***
	(0.0009)	(0.0001)	(0.0001)	(0.0005)
Lagged Dependent Variable (*t* - 1)	0.6468***	0.3614	0.6498***	0.3060**
	(0.1489)	(0.2876)	(0.1878)	(0.1474)
Peak-to-trough effect	-0.11	0.02	0.01	0.06
Peak-to-trough p-value	0.0027	0.0000	0.3437	0.0144
N	134,558	74,776	88,763	59,697
Panel 3. Expenditures ($)				
Unemployment rate	-12.8633***	-2.0324*	-7.1709***	-4.1510***
	(2.3238)	(1.1193)	(1.4857)	(1.2976)
Unemployment rate squared	0.7068***	0.1005	0.3999***	0.2377***
	(0.1477)	(0.0712)	(0.0959)	(0.0825)
Per capita income ($1000s)	0.0561	-0.1363**	0.0126	0.1738***
	(0.0833)	(0.0536)	(0.0659)	(0.0468)
Lagged Dependent Variable (*t* - 1)	0.8484***	0.8159***	0.8237***	0.8644***
	(0.0314)	(0.0430)	(0.0465)	(0.0588)
Peak-to-trough effect	-29.81	-5.15	-16.33	-9.26
Peak-to-trough p-value	0.0000	0.0182	0.0000	0.0003
N	202,362	202,362	202,362	202,362

Notes: Observations are aggregated by household and year. Each panel and column represents results from the preferred linear dynamic panel-data regression model (see text for more details). For all models, the estimation sample excludes households with no alcohol purchases. All models include controls for household demographics, income, employment, area-level characteristics as well as indicators for year and state of residence. Robust standard errors clustered by household are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 9. State-level unemployment and expenditures (total and alcohol)

	Alcohol expenditures ($s)	All expenditures ($1000s)
Unemployment rate	-12.8633***	-0.2314***
	(2.3238)	(0.0493)
Unemployment rate squared	0.7068***	0.0125***
	(0.1477)	(0.0030)
Per capita income ($1000s)	0.0561	0.0010
	(0.0833)	(0.0013)
Lagged dependent variable (*t* - 1)	0.8484***	0.8919***
	(0.0314)	(0.0440)
Peak-to-trough effect	-29.8146	-0.5464
Peak-to-trough p-value	0.0000	0.0000

Notes: N = 202,362. Observations are aggregated by household and year. Each panel and column represents results from the preferred linear dynamic panel-data regression model (see text for more details). For all models, the estimation sample excludes households with no alcohol purchases. All models include controls for household demographics, income, employment, area-level characteristics as well as indicators for year and state of residence. Robust standard errors clustered by household are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table A1. State-level unemployment and annual alcohol purchases (ABV), all coefficients

Unemployment rate	0.4050	-3.1137**	-2.3772***	-0.4821**	-3.3474***	-3.4684***	-3.0785***
	(0.3449)	(1.5240)	(0.8874)	(0.2289)	(1.0935)	(1.0984)	(1.0807)
Unemployment rate squared					0.1822***	0.1802***	0.1553**
					(0.0696)	(0.0691)	(0.0685)
Per capita income ($1000s)	-0.8887***	-1.1350	-0.5534	-0.0933	-0.0788	-0.1413*	0.0391
	(0.1476)	(1.1455)	(0.6349)	(0.0748)	(0.0750)	(0.0788)	(0.0854)
Lagged ABV (*t* - 1)				0.8257***	0.8218***	0.8212***	0.8222***
				(0.0329)	(0.0336)	(0.0336)	(0.0333)
Male head present						16.8019***	16.6357***
						(5.5837)	(5.5517)
Household widowed						6.7172**	6.5489*
						(3.3943)	(3.3754)
Household divorced/separated						10.0924***	9.8926***
						(3.5198)	(3.4969)
Household single						7.8057**	7.5522**
						(3.6247)	(3.6199)
$30k ≤ Household income < $60k						1.6340	1.5216
						(1.4744)	(1.4656)
Household income ≥ $60k						6.6691***	6.3923***
						(1.9204)	(1.8909)
Female employ hrs < 30						-5.9776*	-5.6392*
						(3.3145)	(3.3087)
30 ≤ Female employ hrs < 35						-4.2956	-3.9344
						(4.6228)	(4.6169)
Female employ hrs ≥ 35						-4.0686	-3.7738
						(3.3245)	(3.3188)
Female not employed						-3.3151	-3.0516
						(3.3809)	(3.3790)
Male employ hrs < 30						1.6335	1.6635
						(4.3668)	(4.3672)
Male employ hrs ≥ 35						-3.5872	-3.4368
						(2.8162)	(2.8055)
Male not employed						2.7349	2.7447
						(2.7940)	(2.7885)
Beer tax per gallon							4.0343
							(2.6217)

Beer, wine, & liquor stores							-0.0053
							(0.0116)
→ Employment							-0.0034
							(0.0021)
Drinking places							-0.0096
							(0.0067)
→ Employment							0.0025***
							(0.0008)
Supermarkets							-0.0034
							(0.0034)
→ Employment							0.0003*
							(0.0001)
Convenience stores							-0.0098
							(0.0098)
→ Employment							0.0009
							(0.0021)
Model	OLS	State & Year FE	Household & Year FE	Dynamic PD	Dynamic PD	Dynamic PD	Dynamic PD
Standard Errors	Robust	Clustered (State)	Clustered (Household)	Clustered (Household)	Clustered (Household)	Clustered (Household)	Clustered (Household)

Notes: N = 202,362. Observations are aggregated by household and year. For all models, the dependent variable is ABV (alcohol by volume) purchased. The estimation sample excludes households with no alcohol purchases (16.0%). A household fixed effects model that excluded households who moved between states (2.7%) and that implements standard errors clustered at the state level yields results nearly identical to those reported here (standard errors clustered at the household level).

*** p<0.01, ** p<0.05, * p<0.1