

**Beer in Good Times and Bad: A U.S. State-Level Analysis of
Economic Conditions and Alcohol Consumption**

by

Donald G. Freeman, Ph.D.
Professor & Chair
Department of Economics & International Business
Sam Houston State University
P.O. Box 2118
Huntsville, Texas 77341-2118
tel: 936-294-1264
fax: 936-294-3488
e-mail: freeman@shsu.edu
url: http://www.shsu.edu/~eco_www/faculty/freeman.html

Paper to be presented to the first *Beeronomics Conference*, Leuven, Belgium, May 2009.

DRAFT: For discussion purposes only. Not to be quoted without permission. ©2009.

Beer in Good Times and Bad: A State-Level Analysis of Economic Conditions and Alcohol Consumption

I. Introduction

How economic conditions affect alcohol consumption is still a matter of debate. The convention wisdom in the economics literature is that alcohol is a procyclical normal good, as evidenced by Ruhm (1996), Blake and Nied (1997), Ruhm & Black (2002), and Freeman (1999, 2000). Income elasticities for alcohol tend to vary around 0.5-0.8 in these studies, with per capita consumption of all beverages, including beer, wine, and spirits, falling during recessions. Tremblay and Tremblay (2005), in a comprehensive study of the brewing industry, report on eight studies of beer demand, including their own, six with positive income elasticity and two with negative income elasticity, averaging about 0.2.

The prevailing view from the investment side, however, is that the beer industry is recession-proof or acyclical: beer companies (and other purveyors of “vice” products like tobacco or gambling) are considered to be “defensive” stocks, relatively immune to the business cycle. For instance, over the five year period from June, 1998 to June, 2003, marked by the bursting of the “dot-com” bubble, the recession of 2000-2001, and the tragedy of 11-September-2001, the stocks of companies in the Alcoholic Beverages, Gambling and Casinos, and Tobacco industries rose by 46.02%, 145.18%, and 56.70% respectively, while the S&P 500 lost 14.05% (Ahrens, 2004). Notably, common stocks of brewing companies typically carry “betas” less than 1.0, meaning that their systematic risk is less than the market.¹

Psychological theories suggest on the other hand that alcohol is countercyclical, as economic downturns are causes of stress-induced drinking. In a series of papers Brenner and his

¹ Broadly speaking, “beta” is the coefficient of the return of a security regressed on the return of the market. A beta of 1.0 indicates that the security has the same systematic risk as the market.

colleagues (e.g., Brenner and Mooney, 1983) produced evidence that a host of self- and other-destructive activities, including alcohol abuse and drunk driving, increased during periods of unemployment. Brenner's work was later attacked on methodological grounds, but his work is still cited as influential.

In a recent attempt to measure directly the cyclical nature of beer consumption, Freeman (2001) used monthly shipments data in the U.S. from 1955 through 1994 in an error-correction framework to test for the short-run response to economic factors such as unemployment, industrial production, and personal income. Although there does appear to be a long-run relationship between beer consumption levels and the economic variables, Freeman finds no evidence of short-run cyclical response of beer to economic variables. Freeman's study focuses only on beer, only at the national level, and only using monthly data, but does find supporting evidence of beer's acyclicity.

In order to shed further light on the cyclical nature of beer consumption specifically, the current paper uses U.S. state-level shipments data from 1970 to 2007 as provided by The Beer Institute (formerly the United States Brewers' Association) to estimate pooled time-series models of annual beer consumption on economic and demographic variables. The advantages of using state-level data include the additional variability of regressors, which helps to mitigate problems of multicollinearity common to macro data, the ability to control for state- and time-specific sources of variation, including the different timing and severity of recessions across states, and the greater potential power of tests of stationarity and cointegration.

This paper also incorporates controls for changes in the age distribution of the states, demonstrated by Kerr, et al. (2004) to be an important source of variation in alcohol consumption over time in the U.S. Kerr, et al., using the U.S. National Alcohol Survey, show

that age groups have different consumption habits regarding types and amounts consumed of alcoholic beverages. Tremblay and Tremblay (2005), using market research data, report that the prevalence of beer drinking, while declining somewhat for all groups, is highest (at around 50% in 2001) in the 25-34 year age group. As we will demonstrate below, changes in this group's proportion of the population show a close association with trends in per capita consumption of beer in the U.S. and in the several states.

To estimate these relationships, this paper uses recent developments by Pesaran (2006) and his coauthors to control for unobserved common effects in pooled cross-sections. The traditional way of controlling for common temporal effects like nationwide economic shocks has been to employ a two-way fixed effects model with "dummy" variables for each time period in a pooled estimation. The limitation of the traditional approach is that all the cross-section units are constrained to have the same response to the (single) common time effect. In the Pesaran approach, referred to as common correlated effects (CCE) estimators, the cross-section means of the dependent variable and the regressors are used as proxies for the unobserved common effects in augmented regressions, allowing for heterogeneous responses to the proxies across the units. The CCE estimators are shown to reduce the substantial bias and size distortions that result if cross-section dependence is ignored.

Using the improved techniques, the empirical analysis finds that beer consumption is in fact cyclical, varying negatively with the unemployment level, a result at variance with previous research and one that comes out only when unobserved heterogeneity is properly controlled. Beer is a normal good, varying positively with income, and previous findings that excise taxes have a negative effect on consumption are supported, though not with the same degree of responsiveness. Demographics have a significant and material affect on consumption; as

expected, the larger the share of young adults in the population, the greater the consumption of beer per capita.

The paper proceeds as follows. The next section describes the data and provides a summary of previous research. Section III describes the empirical approach and the initial CCE estimates. Section IV extends these results and provides some robustness tests, and Section V concludes.

II. Previous Research and Proposed Methodology

The main focus of the paper is the responsiveness of beer demand to the business cycle. As measures of economic activity, we use the unemployment rate, the employed/population ratio (for the working-age population), and disposable income per capita, in 2007 dollars, all at the state level, for the years 1970-2007. Each of these variables captures a different characteristic of the business cycle, and the use of multiple regressors is a reflection of the fact that there is no variable that is the one best measure of the cycle. The unemployment rate and the employed/population ratio are taken from the Local area Unemployment Survey of the Bureau of Labor Statistics, and state disposable income is taken from the regional accounts database of the Bureau of Economic Analysis.

Previous research on drinking and the business cycle has produced mixed results, depending on the data level and the econometric specification. As noted, research based on psychological theories, such as Brenner and Mooney (1983), or Winton, et al. (1986), often assert that the stress of unemployment increases drinking behavior as a coping mechanism. Economic research such as Sloan, et al. (1995), or Ruhm (1995) typically regards alcohol as a

normal good, with declining incomes in recessions leading to predictions of declines in alcohol consumption. With no clear theoretical prediction, the question of alcohol's procyclicality becomes an empirical one.

Ruhm (1995) uses a panel of aggregate data to investigate the role of the macroeconomy in relation to alcohol consumption and traffic fatalities, finding that drinking and vehicle mortality decline during recessions. Freeman (1999) looks at an expanded data set corrected for non-stationarity and verifies Ruhm's finding of alcohol's procyclicality. Dee (2001), on the other hand, using micro data from the Behavioral Risk Factor Surveillance System (BRFSS), finds that that binge drinking increases markedly during recessions, and argues that economic stress is the primary factor, outweighing the income effects of recessions. Ruhm and Black (2002) counter with an expanded BRFSS sample with a larger set of explanatory variables and a weighted regression to find again that alcohol is procyclical even at the micro level.

Freeman (2001) takes a somewhat different tack from previous literature to focus solely on beer consumption at the national level, at a monthly frequency, over the period 1955-1994. Because recessions are relatively short-lived phenomena (the average post-WWII recession has lasted only 10 months), the annual frequency used in previous research may not be adequate to pick up responses in alcohol demand to changes in cyclical economic variables. Freeman uses error-correction methodology (ECM) to test the presence of a cointegrating vector among beer consumption, the unemployment rate, personal income and beer excise taxes. Freeman finds evidence of cointegration among the variables, but tests of Granger causality reveal no evidence that beer consumption is responsive to departures from equilibrium, indicating that beer is in fact neither pro nor counter-cyclical.

Most of previous research on alcohol demand using aggregate data uses national data for relatively long time series, or panels of states using relatively short time series. Ruhm (1995), for example, uses a panel of states for 14 years covering only two national business cycles. By contrast, this paper uses 28 years of data for the 50 states plus the District of Columbia, a time period covering four full business cycles and parts of two others. The additional degrees of freedom and the greater variety of economic circumstances allow for a greater precision in estimating the coefficients of interest.

The data used in this paper is taken from “Beer Shipments by State” as reported in *The Brewer’s Almanac*, an annual publication of the Beer Institute, formerly the United States Brewers’ Association (Beer Institute, various years). The data are reported by millions of 31-gallon barrels shipped within each state annually from 1970-2007, and then converted to per capita equivalents.² *The Brewers’ Almanac* is also the source of the state and federal excise tax rates on beer, used as a measure of inter-state price variation.

[Figure 1 about here]

Figure 1 displays annual gallons per capita beer consumption for the two largest states, California and Texas, along with the percentage of the population in each state 20 to 35 years of age. There is much useful information in the figure. First is the general pattern of beer consumption in the two states (more or less representative of most states), wherein beer consumption peaks at about 1980 and declines thereafter. Though not shown in the figure, the 1980 peak marked the end of a long period of growth in consumption beginning at least in the early 1960s in both states.

² Other methods of reporting are used, including per capita consumption above a certain age, or beer in terms of ethanol equivalents. We use per capita consumption for all ages because we incorporate the age distribution in our regressors. We acknowledge the variation in alcohol content in beer due to, for example, the introduction of “light” beers, but the statistical methods employed in the estimation will likely capture much of this variation.

The second piece of information is the gap in consumption between the states, which is constant throughout and even widening in recent years. This gap is a clear indication of heterogeneity in consumption across states, heterogeneity that is unlikely to be fully captured by the regressors in the model. The nature of the observed heterogeneity across states can be seen from the descriptive statistics in Table 1.

[Table 1 about here]

Beer consumption per capita has been highest over the sample period in Nevada (with a young population and a large tourist sector), and lowest in Utah, which has a large Church of Latter-Day Saints population, who do not imbibe for religious reasons. The variation in consumption is quite substantial, with the 20 gallon difference between minimum and maximum consumption equivalent to about 200 12 oz. cans of beer per person per year. On average, there has been a slight decline in beer consumption, as noted for California and Texas in Chart 1.

There is considerable variation in the regressors as well, with maximum unemployment rates more than double minimum states, both over the sample and in 2007 as well. Likewise, income is not distributed equally, with the District of Columbia about twice as rich as the poorest states, in both periods. The employment/population ratio has less volatility and less variation than the other cyclical economic measures, and has trended upward over the sample. Excise taxes, meanwhile, differ by a factor of about three between the highest and lowest-taxed state.

Given mobility, similar cultural heritage, and so forth, the range of young people in states' populations is relatively narrow, but there are nonetheless clear differences. Florida and Maine are relatively old, and the District of Columbia is relatively young (as is Utah among states proper). On average, the percentage of young people in the population has risen, then

fallen over the sample period as the children of the “Baby Boomers,” the large population surge in the U.S. following World War II, have matured.

The observed variation in the regressors notwithstanding, there undoubtedly remains much heterogeneity in beer consumption across states that is unobserved. States have very different laws regarding the distribution and sale of alcoholic beverages, for example, and just as the religious beliefs of the Latter-Day Saints have an effect on beer sales in Utah, so may the religious mix in other states have an effect on consumption there.

Similarly, there are sources of temporal heterogeneity that affect all cross-section units in the sample, such as macroeconomic events like interest rate changes, fiscal policy, etc, or social events like changing attitudes toward certain behaviors like the consumption of alcohol. For example, some of the decline in beer consumption per capita from 1980 may be due to changing attitudes regarding alcohol consumption, cohort influences as generations who grew of age in different eras retire, consolidation of brands, etc.

For heterogeneous panels of cross sections where the number of time periods, T , exceeds the number of cross section units, N , seemingly unrelated regressions (SUR) is a potentially efficient estimation strategy. However, in the present case, $T = 28$ and $N = 51$, so SUR is not feasible.³

We therefore must rely on a method of pooling the data to obtain an estimate of the average response of beer consumption to the explanatory variables in the model. The least square dummy variable (LSDV) model relies on fixed effects for the cross sections and time periods to capture unobserved heterogeneity in the data, and is the workhorse of estimating heterogeneous panels. However, the LSDV model assumes that all heterogeneity can be

³ The sample extends back only until 1970 because age distribution by state is available annually beginning at that year.

captured by simple zero-one variables, and that the response of each cross-section unit to the time effects is identical.⁴ Recently, Pesaran and his co-authors (Pesaran, 2006; Pesaran and Tosetti, 2007) have proposed a more flexible technique, the common correlated effects (CCE) estimator, to control for unobserved heterogeneity in panels of moderate size in T and N . The CCE estimator, which uses cross-section averages of the dependent variable and the regressors to augment regressions of the individual units, is quite flexible and easy to use, and has been shown in Monte Carlo simulations to control for unobserved heterogeneity and cross-section dependence far more completely than the usual two-way fixed effects model. The CCE estimator is described more completely in the following section.

The third piece of information in Figure 1 is the relationship between the proportion of young adults in the population and the per capita consumption of beer. As noted, the highest consuming age group is the young adult category, so it is not surprising that trends in consumption would be related to the size of this group. As a first step, we incorporate the states' age distributions in our analyses by including age in the regression model. This paper is one of the few in the literature on alcohol demand that explicitly accounts for the age distribution of the population.

The relationship between the age distribution and beer consumption may be more subtle than the proportion of a single age group may be able to capture, however. For example, an increase in young adults accompanied by an increase in the very young (and a concomitant decrease in the older adult population) may have very different effects from the same increase in young adults but an opposite increase in the other age groups. Therefore there may be some value in examining the effect of changes in the entire age distribution on beer consumption. This

⁴ Another technique is to use state-specific trends. However, in the present case there is no evidence of linear trends, and the use of quadratic trends would likely overdetermine the model.

we do using a polynomial distributed lag (PDL) technique introduced by Fair and Dominguez as an alternative to the single age group.

III. Empirical Framework and Initial Estimations

The basic model for estimation can be expressed as:⁵

$$y_{it} = \alpha_i' \mathbf{d}_t + \beta_i' \mathbf{x}_{it} + e_{it}, \quad (1.0)$$

where y_{it} is consumption of beer in gallons per capita, \mathbf{d}_t is a vector of observed deterministic effects, \mathbf{x}_{it} is a $k \times 1$ vector of regressors, and the errors, e_{it} , have the multifactor structure

$$e_{it} = \gamma_i' \mathbf{f}_t + \epsilon_{it}, \quad (2.0)$$

with \mathbf{f}_t a vector of unobserved common factors, and ϵ_{it} an idiosyncratic error term uncorrelated with \mathbf{f}_t or \mathbf{x}_{it} . We note that $\mathbf{d}_t = 1$ and $\mathbf{f}_t = 1$, $\gamma_i = \gamma$ and $\beta_i = \beta$ is the traditional two-way fixed effects model. In most cases, it is reasonable to assume that the observed and unobserved factors are related, and this relationship can be expressed as

$$\mathbf{x}_{it} = A_i' \mathbf{d}_t + \Gamma_i' \mathbf{f}_t + v_{it} \quad (3.0)$$

where A_i and Γ_i are factor loading matrices and v_{it} are error terms assumed not to be correlated with the common effects. Stacking (1.0) and (3.0), we obtain

$$\mathbf{z}_{it} = B_i' \mathbf{d}_t + C_i' \mathbf{f}_t + u_{it} \quad (4.0)$$

where $\mathbf{z}_{it} = (y_{it}, \mathbf{x}_{it}')$, $B_i = (\alpha_i \ A_i) \begin{pmatrix} 1 & 0 \\ \beta_i & \mathbf{I}_k \end{pmatrix}$, $C_i = (\gamma_i \ \Gamma_i) \begin{pmatrix} 1 & 0 \\ \beta_i & \mathbf{I}_k \end{pmatrix}$. Assuming suitable rank

conditions on C_i , equation (4.0) can be multiplied by $(C_i \ C_i')^{-1} C_i$ and solved for \mathbf{f}_t to give

$$\mathbf{f}_t = (C_i \ C_i')^{-1} C_i (\mathbf{z}_{it} - B_i' \mathbf{d}_t - u_{it}). \quad (5.0)$$

⁵ This discussion follows closely Pesaran (2006) in which the CCE estimator was first described.

Thus the unobserved common factors can be expressed as linear combinations of the dependent variable, the regressors, and the deterministic effects. Pesaran (2006) shows that asymptotically, any weighted average will be consistent, so simple averages will do.

The estimator described above is named the CCE estimator by Pesaran. It offers two advantages over the traditional two-way fixed effects estimator. First, by admitting the possibility of multiple common but unobserved factors, the CCE estimator offers the potential for consistent estimation of the coefficients of the individual regressors (i.e., the β_i in (1.0) above). Because these coefficients, and not the factor loadings of the unobserved factors, are the items of interest in the current exercise, this method is particularly important. Second, by allowing the coefficients of the f_t to vary across units, the CCE estimator can better accommodate heterogeneous panels where some average of responses is nonetheless required.

To operationalize the CCE estimator, we proceed in two steps. First, we run regressions of each cross section on a constant and on simple averages of the dependent variable and the regressors. Then the residuals from the cross section regressions are pooled and regressed against the x_{it} . The effect is very similar to a pooled time series cross section with two-way fixed effects, except that the variable averages replace the time effects and the procedure allows a separate response for each cross-section unit. In Monte Carlo tests of the CCE estimator in generated samples with unobserved common effects that are correlated with the regressors, Pesaran (2004) and Coakley, Fuertes, and Smith (2004) demonstrate that the CCE estimator is the preferred choice on efficiency grounds when compared to the usual two-way fixed effects and demeaned mean group estimators.

Table 2 presents the results of pooled regressions utilizing various forms of controls for cross-section and intertemporal heterogeneity.

[Table 2 about here]

Model (1) presents the results of a pooled sample of all 51 units with only the regressors as controls and using a common intercept and common slopes. Beer is seen to be strongly countercyclical (each additional percentage point of unemployment results in about a one-quarter increase in beer consumed per person), on the one hand, but consumption also rises when the employment/population ratio rises. These two results are not necessarily inconsistent – the employment/population ratio reflects, in addition to cyclical movements, longer term trends in labor participation – but the two variables do tend to move in opposite directions during the business cycle. Beer in Model (1) is an inferior good, and the greater number of young people in the population, as expressed as the percentage of people ages 20-35, tends to increase beer consumption. Taxes have a strong effect on consumption: each dollar of beer excise tax reduces beer consumption by 1.25 gallons per person annually.

We note, however, that the F -test for common intercepts is soundly rejected in model (1), as might be expected from the data presented in Table 1 showing very different levels of beer consumption across the states. In model (2), we introduce fixed state effects into the estimation to control for levels differences across states in beer consumption due to climate, religion, tourism, etc. The fixed effects obviously provide a better fit for the model, as R^2 jumps from about 20 to about 90 per cent. Some of the coefficient on the regressors change significantly as well, with unemployment now having no effect on consumption and the effect of the employment ratio much reduced. The tax effect is almost doubled, however, and the age effect is marginally stronger. There appears to be much remaining cross correlation in the sample, however, as the Pesaran “ CD ” test, specified as $CD = \sqrt{\frac{2T}{N(N-1)} (\sum_{i=1}^{N-1} \sum_{j=i+1}^N \hat{\rho}_{ij})}$, where $\hat{\rho}_{ij}$ are the contemporaneous cross correlations of the disturbances and distributed as a standard normal

variant, soundly rejects the null of zero correlation, and the pooled Durbin-Watson indicates the presence of within-unit serial correlation as well.

As noted earlier, the age distribution may affect beer consumption in more complex ways than a single age group can capture. The rationale for including the entire distribution is that the proportions of ages older or younger than normal beer-drinking age may modify the drinking of the latter. For example, the drinking habits of the twenty per cent of the population aged 20-35 may be different if the split in the older and younger population is an even 40-40 instead of say, 50-30 or 30-50.

To proceed, we use Fair and Dominguez's (1991) polynomial distributed lag (PDL) technique as applied to the entire age distribution. In the Fair-Dominguez framework, the age distribution enters the model originally as:

$$y_{it} = \alpha'_t \mathbf{d}_t + \beta'_i x_{it} + \delta_1 p_{1it} + \dots + \delta_n p_{nit} + e_{it} \quad (6.0)$$

with p_{jit} the proportion of age group j in the population, where each j consists of 5-year age groups up to age 65, and then all remaining years. It is clearly not feasible to estimate 14 age group coefficients; besides the degrees of freedom limitations, the age groups are highly collinear, leading to imprecise estimates of the true pattern of the coefficients. Fair and Dominguez therefore propose to impose two restrictions on the coefficients: 1) that they all sum to zero; and 2) that they lie on a second-degree polynomial $\delta_j = a_0 + a_1 j + a_2 j^2$. The zero-sum constraint implies that for n groups $a_0 = -a_1(1/n) \sum_{j=1}^n j - a_2(1/n) \sum_{j=1}^n j^2$, and given the standard formulae for the summation of the first n numbers and first n squares, the age variables will enter (6.0) as $a_1 Z_{1it} + a_2 Z_{2it}$ where $Z_{1it} = \sum_{j=1}^{14} j p_{jit} - 7.5$ and $Z_{2it} = \sum_{j=1}^{14} j^2 p_{jit} - 72.5$.

The results of the Fair-Dominguez PDL treatment of age are shown in column (3) of Table 2. Using the entire age distribution alters the results in a material way for several of the estimated coefficients: the coefficient on unemployment is about twice times as large and is now statistically significant, the coefficient of the employment ratio is less than half the size of model (2), and the income coefficient is also twice the size of Model (2). The summary statistics are little changed, however. The implied shape of the responses of the age distribution are described and illustrated in connection with the discussion of Model (6), below.

Model (4) introduces time fixed effects to further correct for unobserved variation in the sample due to common influences occurring in the time dimension; this is the two-way fixed effects model. Unemployment is now insignificant and the coefficients on income and employment resemble those of Model (2). The coefficients on the age distribution increase, as does the coefficient on taxes, with each additional dollar of excise taxes now causing a more than two and one-half gallon reduction per person reduction in annual beer consumption. Moreover, the CD statistic reduced by a factor of almost four, if still statistically significant. There is no improvement in the Durbin-Watson, however, indicating that fixed time effects are insufficient to account for serial correlation in the residuals.

Model (5) introduces the main innovation in this paper, the use of the CCE estimator. We replace the time fixed effects with a two-step procedure first regressing each cross-section unit on an intercept and the cross-section averages of the set $(y_{it}, \mathbf{x}_{it})$, then pooling and regressing the residuals e_{it}^y on the e_{it}^x . We find that the CCE estimator produces significant changes in the responses of beer consumption to the regressors. Beer consumption is now definitely procyclical, with increases in unemployment producing a small yet significant reduction in beer

consumption. The employment ratio now produces no response in consumption, but the sign of the income variable is also reversed and significant, indicating that beer is a normal good. Of special note is that the coefficient of the tax variable is much reduced. We suspect that the construction of the tax variable by deflating a fixed nominal tax rate with the consumer price index, a common practice in studies of this type, is capturing some of the downward trend in beer consumption in Models 1-3. The cross-section average of the tax variable in the first stage regressions of Model (5) is capturing more or less the same movement, thus the residuals used in the second stepped have been in effect de-trended and the tax variable is no longer significant.

The coefficients of the age variables are much increased in magnitude, indicating a much steeper rise and fall in the effects on beer consumption as we move through the age distribution. The R^2 of Model (5) (calculated using the variance of the original dependent variable) is somewhat improved, and the CD statistic is smaller but still large enough to reject the null of no cross-correlation. A notable difference is the size of the Durbin-Watson statistic, still indicating the presence of serial correlation, but much improved from the earlier models. We can thus be reasonably assured that much of the unobserved heterogeneity is being captured by the CCE model in a way that the traditional two-way LSDV model could not handle.

In Model (6) we add the lagged dependent variable to the CCE model. The coefficients of the regressors are little changed, except for the tax coefficient, but the pooled Durbin-Watson statistic now indicates that the null of no serial correlation cannot be rejected. The consequences of using a lagged dependent variable (LDV) in the CCE model are not well established (although it is certainly known that the use of a LDV in panel data introduces bias, especially in the coefficient of the LDV), but the lack of substantial change in the coefficients indicates that they

must be relatively benign in the current circumstances. The *CD* statistic is also reduced, though still significant.

The estimated coefficients of the age distribution variables from Model (6) are used to generate the δ 's in equation (6.0) so that the effect of the entire age distribution can be displayed, as in Figure 2.

[Figure 2 about here]

In the figure, the markers on the line indicate implied coefficients for the five-year interval ending at that marker. The pattern shows that consumption peaks at age 40, a little later than expected, but not much at variance with the survey results reported by Tremblay and Tremblay (2005). Greater concentration of the population at very young (especially) and very old age groups subtracts from per capita beer consumption, as expected.

Model (6) yields the strongest results in terms of the economic variables; that is, the coefficients of the regressors tell a nice story of beer as a pro-cyclical, normal good that is sensitive to excise taxes and therefore price. The implied price elasticity of beer as derived from the tax coefficient is quite small, however; using the means of the variables from Table 1 results in price elasticity of only -0.045; the income elasticity is also small at +0.041. These estimates are well below those of recent research, including Selvanathan and Selvanathan (2004), Clements and Johnson (1983), and Clements and Selvanathan (1987). Evidently, the results of prior research may have been overstated by the omission of controls for common effects.

IV Extensions and Robustness tests

One criticism of the CCE estimator may be that it groups all states together to derive cross-section means used to proxy the unobserved common effects, thereby ignoring regional

differences in these common effects. Regions may differ in industrial composition, obviously in weather, religious affiliation, etc. In this section we modify the two-step CCE by using regional averages rather than national averages in the first-stage regression to account for unobserved effects. Table 3 provides the results of this modification, along with results from one other extension, to be discussed below.

[Table 3 about here]

Model (7) uses regional rather than national averages, as in Model (6), reproduced from Table 2. Using the regional effects changes the results only slightly, with the net effect that most coefficients are reduced in magnitude. Most likely, the regional averages are more closely correlated to the individual state regressors, so the consequent multicollinearity reduces the size and significance of the coefficients of the regressors. The summary statistics suggest that there is little to gain from using the regional as opposed to the national averages.

In a second extension, we split the sample in half at year 1988 and re-estimate model (5) from Table 2 (the lagged dependent variable is omitted to preserve degrees of freedom). There is no statistical reason for choosing 1988 as the break year; the purpose is to compare the stability of the models using a break, so the choice of year is arbitrary. We conduct the break test for the 2-way fixed effects Model (4) from Table 2, and for the CCE model using variable averages, Model (5) in Table 2.

The results from breaking the sample are not encouraging for a consistent story using the 2-way fixed effects model in columns (8) and (9). In this case, the unemployment coefficient is strongly positive in the second period, and the age coefficients are small and insignificant in the first period. The employment ratio and excise taxes are consistent across the two periods, but

income only in the first. It is obvious that the findings of the fixed effects model are sensitive to the time periods chosen for estimation.

By contrast, the CCE model in (10) and (11) appears to tell a consistent story that beer is procyclical in both halves of the sample. The coefficient of the income variable is significant only in the second half. The employment ratio is never significant, but the age variables are consistent in sign and significance, if not in magnitude. The CCE also models get a slight nod in terms of goodness-of-fit and in tests of serial correlation as compared to the 2-way fixed effects models. The results are far from conclusive, but the CCE model is more robust to arbitrarily specified sample periods.⁶

Other beverages

A potential confounding factor in the estimations above is the changing tastes of consumers for different types of alcoholic beverages. Over time, whether due to income, status concern, cohort effects, or other factors, drinkers may move among beer, wine, or spirits. For example, per capita consumption of spirits in the US had been trending down for many years, reaching a low of 1.21 gallons in 1998, before rebounding by 20% to 1.45 gallons in 2007. By contrast, beer consumption has stayed flat at about 21.8 gallons over the past decade.

A preferred way to control for other beverages would be to estimate a systems model using all three types of beverages in a system of equations similar to Selvanathan and Selvanathan (2004) for national consumption in Australia. Unfortunately, we lack state-level pricing data for any alcoholic beverages, and simply including the consumption of other beverages on the right-hand side of equation (1.0) will likely lead to issues of endogeneity bias,

⁶ A Chow test ($F = 0.61$) fails to reject the null hypothesis of coefficient homogeneity for the CCE model. Of course, a more satisfactory way to test for stability is a search for an endogenous break, but that is beyond the scope of this paper.

as shocks to beer consumption are likely correlated with shocks to consumption of other alcoholic beverages.

The two-step CCE estimator provides a way of controlling for changing national tastes in alcoholic beverages, however, by incorporating wine and spirits in the initial step regressing state beer consumption per capita on the cross-section means of all variables. Thus the initial step will take the form:

$$y_{it} = c_i' \bar{z}_t + \delta_{iw} \bar{w}_t + \delta_{is} \bar{s}_t + \mu_{it} \quad (7.0)$$

where $\bar{z}_t = (\bar{y}_t, \bar{x}_t)$, \bar{w}_t is the cross-section mean of wine consumption per capita, and \bar{s}_t the cross-section mean of spirits consumption per capita. The second step will be as above, the regression of the residuals from (7.0) on the residuals of the regressors from similar first-step regressions.

[Table 4 about here]

The results of the test using wine and spirits consumption are shown in Table 4. The inclusion of the wine cross-section mean alone has very little effects on the results; see model (6) in Table 3. The inclusion of the spirit cross-section mean, however, alone or in combination with wine, has a more pronounced effect on the results, especially regarding the lagged dependent variable, which greatly reduced in magnitude in all cases where spirits are present and statistically insignificant when entering the regression separately. The coefficients on the age distribution variables are also reduced, and in two cases the tax coefficient is rendered insignificant.

In summary, the CCE estimator provides consistent results in a variety of specifications. Cross-section correlation is never completely eliminated, but is reduced substantially from

models controlling only for heterogeneous intercepts. Interestingly, the two-way fixed effects model performs just about as well in tests of cross-correlation, but the economic implications of the coefficients are not consistent with prior research, and in any case do not do well under conditions of an arbitrary break in the sample. The presence of other alcoholic beverages or the use of regional averages do not seem to materially affect the results.

V Conclusion

This paper uses recent developments in the econometrics of pooled time series with unobserved common effects to re-estimate the effects of economic and demographic variables on beer consumption using U.S. state-level data over the period 1970-2007. The principal findings of the paper are that beer is a pro-cyclical, normal good that has been quite responsive to changes in the age distribution of the states during the sample period. These results are in contrast to those from a standard two-way fixed effects model using the same data. The latter model estimates beer consumption as being a non-cyclical, inferior good with much smaller response to the age distribution variables.

Another area where the models differ is in the estimated effects of excise taxes on beer consumption within states across time. The two-way fixed effects models estimate the coefficient on taxes to be about four times as large as the estimate for the CCE estimator, indicating that the effect of taxes may have been overstated in previous research, a finding also reported by Dee (1999) and Young and Likens (2000) in a different context. While we have no evidence on the effect of taxes on, say, heavy versus light drinkers, we can say that taxes have only small effects on aggregate consumption.

The results of the Fair-Dominguez PDL technique of measuring the effects of the age distribution on beer consumptions confirm that increases in the proportion of the very young or very old in the age distribution will result in less beer consumed. The peak of the age distribution occurs in the 35-40 year age category, a bit later than the survey reports of beer drinking by age would indicate. It is difficult to interpret the peak of the distribution as being the peak of consumption, however, given that the entire distribution is being modeled instead of the age group that happens to provide singly the largest coefficient. It is interesting that here, too, the two-way fixed effects model differs, placing the peak even later, at the 40-45 year age category.

Further evidence of the superiority of the CCE estimator over the two-way fixed effects estimator is found in the reduction of both cross-section dependence and especially serial correlation of the residuals using the CCE estimator. Though neither is reduced to zero, the use of cross-section means shows great promise in modeling unobserved common effects.

References

- Ahrens, D (2004). *Investing in Vice: The Recession-Proof Portfolio of Booze, Bets, Bombs, and Butts*. New York, NY: St Martin's Press.
- The Beer Institute (formerly the United States Brewers Association). *The Brewers Almanac* (various issues). Washington, DC: The Beer Institute.
- Blake, D and A. Nied. (1997). "The Demand for Alcohol in the United Kingdom," *Applied Economics* 29, pp. 1655-1672.
- Brenner, MH and A Mooney (1983). "Unemployment and Health in the Context of Economic Change," *Social Science & Medicine* 17, pp 1125-1138.
- Clements, KW and LW Johnson (1983) The Demand for Beer, Wine and Spirits: a System-wide Analysis," *Journal of Business* 56 (3), pp 273-304.
- Clements, KW and EA Selvanathan (1987) "Alcohol Consumption," in *Applied Demand Analysis: Results from System-Wide Approaches*, H. Theil and KW Clements, eds. Cambridge, MA: Ballinger Publishing Co.
- Coakley, J, A Fuertes and R Smith (2004). "Unobserved Heterogeneity in Panel Time Series Models," Working Paper BWPEF 0403, University of London Birbeck.
- Dee, TS (1999). "State Alcohol Policies, Teen Drinking, and Traffic Fatalities," *Journal of Public Economics* 72, pp 289-315.
- Dee, TS (2001). "Alcohol Abuse and Economics Conditions: Evidence from Repeated Cross-Sections of Individual-Level Data," *Health Economics* 10, pp 257-270.
- Fair, RC and KM Dominguez (1991). "Effects of the Changing US Age Distribution on Macroeconomic Equations," *American Economic Review* 81 (5) pp 1276-1294.

- Freeman, DG. (1999). "A Note on Economics Conditions and Alcohol Problems," *Journal of Health Economics* 18 (5), pp. 661-670.
- Freeman, DG. (2000). "Alternative Panel Estimates of Alcohol Demand, Taxation, and the Business Cycle," *Southern Economic Journal* 67 (2), pp. 325-344.
- Freeman, DG. (2001). "Beer and the Business Cycle," *Applied Economics Letters* 2001 (8), pp. 51-54.
- Kerr, WC., TK Greenfield, J Bond, Y Ye, and J Rehm (2004). "Age, Period, and Cohort Influences on Beer, Wine, and Spirits Consumption Trends in the US National Alcohol Surveys," *Addiction* 99, pp. 1111-1120.
- Pesaran, MH (2004). "General Diagnostic Tests for Cross Section Dependence in Panels," *CEsifo Working Paper series* n. 1229.
- Pesaran, M. H. (2006), 'Estimation and inference in large heterogeneous panels with a multifactor error structure'. *Econometrica* 74, 967-1012.
- Pesaran, MH and E Tosetti (2007). "Large Panels with Common Factors and Spatial Correlations," *CEsifo Working Paper series* n. 2103
- Ruhm, CJ. (1996). "Economic Conditions and Alcohol Problems," *Journal of Health Economics* 14, pp. 583-603.
- Ruhm, CJ, and WE Black (2002). "Does Drinking Really Decrease in Bad Times?" *Journal of Health Economics* 21, pp. 659-678.
- Selvanathan, EA, and S Selvanathan (2004). "Economic and Demographic Factors in Australian Alcohol Demand," *Applied Economics* 36, pp. 2405-2417.
- Sloan, FA, BA Reilly and C Schenzler (1995). "Effects of Tort Liability and Insurance on Heavy Drinking and Drinking and Driving," *Journal of Law and Economics* 38, pp 49-78.

Tremblay, VJ and CH Tremblay (2005). *The US Brewing Industry: Data and Economic Analysis*. Cambridge, MA: The MIT Press.

Winton, M, N Heather, and I Robertson (1986) “Effects of Unemployment on Drinking Behavior: a Review of the Relevant Evidence,” *The International Journal of Addictions* 21, pp 1261-1283.

Young, DJ and TW Likens (2000) “Alcohol Regulation and Auto Fatalities,” *International Review of Law and Economics* 20, pp 107-126.

Table 1: Descriptive Statistics, Beer Consumption and Explanatory Variables, annual data, 1970-2007, fifty states and the District of Columbia

Variable		Minimum	Maximum	Average	Std. Deviation
Beer, Gallons <i>per capita</i>	1970-2007	13.1 (Utah)	33.5 (Nevada)	22.9	3.7
	2007 only	12.0 (Utah)	32.4 (North Dakota)	22.8	3.9
Unemployment Rate	1970-2007	3.4 (Nebraska)	8.3 (West Virginia)	5.8	1.1
	2007 only	2.6 (Hawaii)	7.2 (Michigan)	4.4	1.0
Employment/Population	1970-2007	48.7 (W. Virginia)	67.5 (Minnesota)	61.5	3.7
	2007 only	53.4 (West Virginia)	71.2 (N. Dakota)	64.1	3.9
Disposable Income, 2007 \$ ('000)	1970-2007	12.6 (Mississippi)	23.1 (D.C.)	16.7	2.3
	2007 only	18.3 (Utah)	38.6 (D.C.)	23.6	4.0
Excise Taxes. 2007 \$/gallon	1970-2007	1.23 (Wyoming)	3.65 (S. Carolina)	1.85	0.60
	2007 only	0.97 (Wyoming)	2.90 (Alaska)	1.38	0.40
% Population ages 20-35	1970-2007	14.1 (Florida)	19.6 (D.C.)	15.7	0.95
	2007 only	11.7 (Maine)	18.2 (D.C.)	13.9	1.28

Table 2: Pooled Estimations of Per Capita Beer Consumption, Annual U.S. States, 1970-2007 (*p*-values in parentheses)

	Model					
Variable	(1)	(2)	(3)	(4)	(5)	(6)
Lagged Dependent Variable	-	-	-	-	-	0.245 (.000)
Unemployment	0.275 (.000)	0.029 (.267)	0.050 (.092)	0.016 (.680)	-0.141 (.000)	-0.139 (.000)
Employment/Population	0.284 (.000)	0.134 (.000)	0.047 (.079)	0.110 (.000)	0.002 (.955)	-0.003 (.806)
Real Disposable Income	-0.057 (.046)	-0.104 (.000)	-0.262 (.000)	-0.100 (.002)	0.083 (.067)	0.075 (.094)
Excise Taxes	-1.252 (.000)	-2.391 (.000)	-2.544 (.000)	-2.649 (.000)	-0.731 (.003)	-0.467 (.052)
Age 20-35	0.261 (.000)	0.331 (.000)	-	-	-	-
Z_1 Age PDL	-	-	12.28 (.000)	15.60 (.000)	42.28 (.000)	34.94 (.000)
Z_2 Age PDL	-	-	-0.772 (.000)	-0.851 (.000)	-2.661 (.000)	--2.181 (.000)
Fixed State Effects	No	Yes	Yes	Yes	Yes	Yes
Fixed Time Effects	No	No	No	Yes	No	No
Variable Cross Section Means (CCE)	No	No	No	No	Yes	Yes
<i>R-Square</i>	0.198	0.879	0.865	0.883	0.982	0.983
Pesaran “ <i>CD</i> ” Heterogeneity Test	13.26 (.000)	23.60 (.000)	22.71 (.000)	5.90 (.000)	5.02 (.000)	4.46 (.000)
Durbin-Watson	0.042 (.000)	0.289 (.000)	0.265 (.000)	0.245 (.000)	1.490 (.000)	1.955 (.212)
<i>F-test for $\alpha_i = \alpha$</i>	211.5 (.000)	-	-	-	-	-


Table 3: Extensions of the CCE Estimators of Per Capita Beer Consumption,
Annual U.S. States, 1970-2007 (*p*-values in parentheses)

Variable	Model					
	National	Regional	2-way fixed effects		CCE estimator	
	(6)	(7)	(8)	(9)	(10)	(11)
		1970-88	1989-2007	1970-88	1989-2007	
Lagged Dependent Variable	0.279 (.000)	0.223 (.000)	-	-	-	
Unemployment	-0.086 (.001)	-0.093 (.001)	-0.013 (.734)	0.269 (.000)	-0.152 (.001)	-0.099 (.018)
Employment/Population	-0.013 (.579)	0.002 (.965)	0.101 (.009)	0.154 (.001)	0.023 (.594)	-0.018 (.619)
Real Disposable Income	0.200 (.000)	0.107 (.020)	0.148 (.063)	-0.032 (.500)	-0.096 (.202)	0.234 (.000)
Excise Taxes	-0.455 (.054)	-0.242 (.262)	-1.445 (.000)	-1.872 (.000)	-0.582 (.142)	-0.874 (.004)
Z ₁ Age PDL	32.51 (.000)	10.53 (.007)	2.468 (.569)	29.95 (.000)	50.27 (.000)	36.41 (.000)
Z ₂ Age PDL	-2.021 (.000)	-.424 (.110)	0.068 (.813)	-1.592 (.000)	-3.239 (.000)	-2.240 (.000)
Fixed State Effects	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Time Effects	No	No	Yes	Yes	No	No
Variable Cross Section Means (CCE)	Yes	Yes	No	No	Yes	Yes
<i>R-Square</i>	0.982	0.986	0.928	0.929	0.978	0.979
Pesaran “ <i>CD</i> ”	4.39	4.74	5.37	4.42	3.89	5.74
Heterogeneity Test	(.000)	(.000)	(.000)	(.000)	(.000)	(.000)
Durbin-Watson	1.945 (.151)	1.959 (.245)	0.554 (.000)	0.317 (.000)	1.693 (.000)	1.187 (.000)

Table 4. CCE Estimators with cross-section means of Wine and Spirits Consumption, annual data, 1970-2007 (*p*-values in parentheses)


	Other Beverage Included			
	Wine Only (12)	Spirits Only (13)	Wine + Spirits (14)	Wine, Spirits (15)
Lagged Dependent Variable	0.235 (.000)	0.036 (.111)	0.048 (.000)	-0.001 (.957)
Unemployment	-0.131 (.000)	-0.116 (.000)	-0.129 (.000)	-0.090 (.003)
Employment/Population	-0.006 (.807)	0.0123 (.637)	0.015 (.577)	0.002 (.941)
Real Disposable Income	0.069 (.123)	0.100 (0.032)	0.081 (.087)	0.105 (0.026)
Excise Taxes	-0.632 (.008)	-0.398 (.156)	-0.359 (.176)	-0.737 (.011)
Z_1 Age PDL	35.54 (.000)	18.54 (.001)	27.81 (.000)	14.81 (.001)
Z_2 Age PDL	-2.241 (.000)	-1.028 (.000)	-1.719 (.000)	-0.863 (.003)
Fixed State Effects	Yes	Yes	Yes	Yes
Fixed Time Effects	No	No	No	No
Variable Cross Section Means (CCE)	Yes	Yes	Yes	Yes
<i>R-Square</i>	0.984	0.987	0.986	0.987
Pesaran “ <i>CD</i> ” Heterogeneity Test	4.48 (.000)	4.05 (.000)	4.15 (.000)	4.12 (.000)
Durbin-Watson	1.993 (.508)	1.982 (.681)	1.971 (.318)	2.014 (.681)

Figure 1



Note to Figure 1: The line with long dashes is annual gallons per-capita beer consumption in Texas and the line with short dashes is the same for California; the line with short and long dashes is the percent of the population ages 20 to 35 in Texas and the solid line the same for California.

Figure 2



Note to Figure 2: Taken from Model (6), Table 2. Implied Coefficients for five-year age intervals from a polynomial distributed lag (PDL) estimation with summation constraint zero. See Fair and Dominguez (1991).