IS BEER CONSUMPTION IN IRELAND ACYCLICAL?

GEARÓID GIBBS
Senior Sophister

In this econometric investigation, Gearóid Gibbs examines beer consumption in Ireland and its relation to the business cycle. Citing psychological studies on increased alcohol consumption during recessions, he hypothesises that this relationship may be an acyclical one, while also emphasising evidence for procyclical consumption from economic literature. While the results of the time-series regression analysis were inconclusive, the scope for further research in this area is emphasised.

Introduction

Ireland's reputation as a nation of prolific drinkers is well established. Of all alcoholic drinks, beer continuously tops the list and accounted for 48 per cent of total alcohol consumption in Ireland in 2010, equating to approximately 98 litres of beer or over 200 pints per person (WHO, 2010). In light of these statistics, this paper seeks to examine the nature of the consumption of beer in Ireland.

Traditionally, investors have viewed alcoholic beverages, particularly beer, as products which are relatively unaffected by trends in the business cycle. Indeed, some identify drink manufacturers' stocks as being 'defensive' against downturns. Empirical research shows the returns on the common stocks of big beer firms are about 50 per cent less volatile than the market average (Freeman, 2001).

This phenomenon raises questions regarding the economic nature of beer. Is beer a normal good, with demand increasing as income increases? Or is beer an inferior good, resulting in decreased demand as income increases? Psychology also has relevant inputs with some theories indicating that alcohol may help to alleviate economic anxieties. This suggests an acyclical relationship between alcohol consumption and the business cycle. Given beer is a relatively cheap beverage, as compared to wine or spirits, individuals may substitute more expensive drinks for beer. In this case, there is a two-fold force behind increasing beer consumption in a recession. This paper aims to investigate the cyclicality of beer consumption in Ireland, with the specific hypothesis that beer consumption is acyclical.

The paper is outlined as follows: section two reviews the background and the existing literature, sections three and four specify the empirical approach and data used

in the study, section five presents the empirical results, and section six discusses possible extensions to the research.

Background and Literature Review

The hypothesis of acyclical beer consumption is developed from several psychological studies that have shown individuals to increase their alcohol consumption during economic downturns. Brenner and Mooney (1983) contend that individuals suffering from the 'stresses of unemployment' may attempt to 'alleviate psychological distress by medication with alcohol'. The authors found that as unemployment rates rise, self-destructive activities such as alcohol abuse and drunk driving increase in prevalence.

The economic literature is more varied. The general view is that alcohol is a pro-cyclical normal good, with consumption rising in income. Tremblay and Tremblay (2005) conducted an analysis of the US Brewing Industry. In a summation of eight previous studies on beer demand, they show six results of beer being a normal good, and two findings of an inferior good. Freeman (2001) analysed beer and the business cycle in the US over the period January 1955 to December 1994. Using monthly data, his estimation results in a co-integrating relationship between beer, beer taxes and cyclical economic factors suggesting that beer is somewhat immune to economic cycles. 'Beer consumption increases with income and industrial production, but increases in unemployment also'. Consistent with this finding of beer consumption being acyclical, Blake and Nied (1997) conclude that the consumption of beer in the UK, in contrast with three other types of alcohol beverage, increases with the unemployment rate. However, they also find the long run income elasticity of beer to be 0.8 suggesting that beer consumption is still somewhat affected by cyclical economic factors.

Bor *et al.* (2013) found that during the 'Great Recession' abstinence rates from alcohol increased in the US. However, this decrease was countered by an increase in total alcohol consumption rates, as more people became 'frequent binge drinkers'. 'The rise in frequent binging was observed for both employed and unemployed respondents, suggesting that factors other than job loss were driving these changes'. The authors contend that this diverging result can be explained by the countervailing 'income-effect' hypothesis and 'provocation' hypothesis. On the one hand, lower income reduces consumption, especially amongst low-income groups; however, recession-linked insecurities may lead to greater alcohol intake as a stress alleviation mechanism.

The authors' conjecture that the downturn associated with the Great Recession was much more severe than the previous business-cycle fluctuations that have been analysed in the literature. This result of countercyclical binge drinking is confirmed by Dee (2001), who found that a five percentage point increase in the unemployment rate results in an 8 per cent increase in the probability of heavy drinking.

Empirical Approach

A time-series regression model of annual beer consumption on cyclical economic and demographic variables is estimated. The model is specified as:

$$Yi = ß0+ß1X1+ ß2X2+ ß3X3+ ß4X4+ ß5X5+ ß6X6 \quad (1)$$

where:
Yi: BEER - the annual recorded consumption of pure alcohol (corresponding to beer at 5%) in litres, per person, aged 15 years old and over.
X1: UNEMPLOYMENT - the seasonally adjusted annual average standardised unemployment rate.
X2: LOGGDPPC - the log of Gross Domestic Product per capita.
X3: EMPLOYPOP - the ratio of those in employment to the population.
X4: PROPYOUTH - the proportion of the population aged between 25-34.
X5: EXCISE - the amount of excise tax on a pint of beer in euro.
X6: VAT - the amount of VAT on a pint of beer in euro.

Data and Expectations

The data is aggregated from several sources and covers the period from 1983 to 2013. Summary statistics are presented in Table 1. The dependent variable is the annual recorded consumption of pure alcohol (corresponding to beer at 5 per cent) in litres, per person, aged 15 years old and over. This is obtained from the World Health Organisation. Figure 1 shows the development of beer consumption in Ireland over the past three decades.

On the right hand side of the regression equation several cyclical variables are included. These are the unemployment rate, the seasonally adjusted annual average standardised unemployment rates taken from the Central Statistics Office; employed/population ratio for the population aged 15 years and older; and, Gross Domestic Product per capita, obtained from the World Bank statistics database. Different variables are included to capture varying trends in the business cycle. If the hypothesis of acyclical consumption is correct, we would expect a positive relationship between beer consumption and unemployment. Increasing unemployment should increase the consumption of beer. In a recession the employed/population ratio is likely to fall, we would then expect a negative relationship between this variable and the consumption of beer. GDP per capita is a measure of individual income. We would expect beer consumption to increase as GDP decreases, but if beer is a normal good; its consumption would increase as income increases. The expected effect is ambiguous and depends on the relevant magnitude of each effect.

Given the distribution of per capita GDP is badly skewed, a non-linear relationship is created between GDP per capita and beer consumption. To control the skew GDP per capita is transformed by taking its logarithm, this action can also help to prevent het-

eroskedasticity problems.

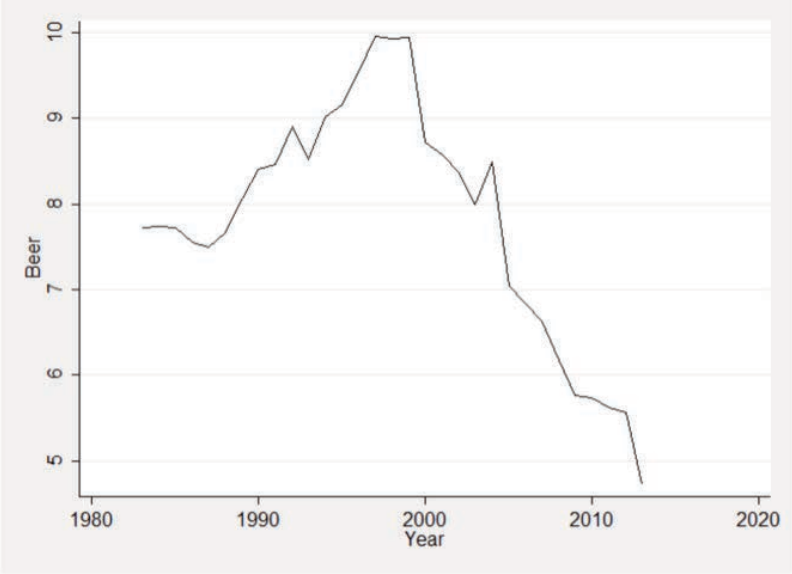

Figure 1: Beer Consumption (litres of alcohol) from 1983 to 2013

Kerr *et al.* (2004) document that different age groups have varying consumption patterns regarding types and quantities of alcoholic drinks consumed. Tremblay and Tremblay (2005) show for the US that individuals in the 25-34 year cohort have the highest consumption of beer (approximately 50 per cent in 2001). Following Freeman (2011), we control for changes in the size of this group. As this is the highest consuming age group, we would expect trends in consumption to be related to the size of this group.

Two tax variables are included. Tax on beer includes excise and VAT. Data comes from the Revenue Commissioners who provide statistics on the tax liability levied on a pint of beer. Given that beer may be addictive, government may attempt to reduce consumption by increasing the price. Price increases are generally borne by increasing the relevant excise rate. Tax may therefore have been used as a policy measure to attempt to reduce alcohol consumption and its related health and social harms. We would expect a negative relationship between tax and beer consumption.

There are several issues and assumptions about the data on beer that must be recognised. Firstly, statistics are only as accurate as data recorded. When considering alcohol consumption, it should be recognised that some consumption may not be recorded.

Alcohol that is homemade, smuggled, or intended for purposes other than consumption may go unrecorded. The World Health Organisation estimated Irish unrecorded alcohol consumption to be approximately 0.5 litres per head (for the population aged 15+ years) in 2010.

Variable	Obs	Mean	Std. Dev.	Min	Max
Beer Consumption	31	7.809	1.394	4.73	9.96
Unemployment	31	10.974	4.803	3.9	17
Log GDP per capita	31	10.195	0.389	9.557	10.659
Employment to Population	31	50.83	5.857	43.8	61
Excise	31	0.445	0.033	0.37	0.47
VAT	31	0.501	0.148	0.34	0.75
25-34 population proportion	31	0.151	0.012	0.772	1.19

Table 1: Summary Statistics for Variables

There are several issues and assumptions about the data on beer that must be recognised. Firstly, statistics are only as accurate as data recorded. When considering alcohol consumption, it should be recognised that some consumption may not be recorded. Alcohol that is homemade, smuggled, or intended for purposes other than consumption may go unrecorded. The World Health Organisation estimated Irish unrecorded alcohol consumption to be approximately 0.5 litres per head (for the population aged 15+ years) in 2010.

Secondly, given that the volume of beer consumed is based on the Revenue Commissioners clearances data, we assume that all beer available for consumption in a particular year is consumed in that year. This does not account for beer that may be imported and subsequently re-exported, beer that has been stored, or beer used in the preparation of food or discarded.

Empirical Results

The first regression model is estimated. Table 2 outlines the preliminary results from the time-series OLS regression. The R-squared of 0.9236 suggests the model explains 92.36 per cent of the variation in beer consumption. This seems rather high for a very simple model.

Unemployment, employment to population ratio, and the 25 to 35 population proportion are insignificant. Log of GDP per capita, excise, and vat are all statistically significant at 5 per cent. Before drawing inference from these results it is necessary to check the model for potential problems. Starting with heteroskedasticity, the square of the resid-

uals is plotted against the fitted values in Figure 2.

Beer	Coefficient	Std. Err.	t	p-value
Unemployment	-0.348	0.168	-0.21	0.838
Log GDP per capita	4.582	1.027	4.46	0.000
Employment to Population	-0.087	0.183	-0.48	0.637
Excise	13.527	4.576	2.96	0.007
VAT	-10.762	3.142	-3.43	0.002
25-34 population proportion	-49.743	38.989	-1.28	0.214
Constant	-27.187	13.995	-1.94	0.064

Table 2: OLS Regression Output

Unemployment, employment to population ratio, and the 25 to 35 population proportion are insignificant. Log of GDP per capita, excise, and vat are all statistically significant at 5 per cent. Before drawing inference from these results it is necessary to check the model for potential problems. Starting with heteroskedasticity, the square of the residuals is plotted against the fitted values in Figure 2.

The plot of the residuals shows a relatively even distribution across the fitted values. A more formal test for heteroskedasticity is also performed. The Breusch-Pagan/Cook Weisberg test is designed to detect linear forms of heteroskedasticity. A small chi-square value confirms that heteroskedasticity will not be an issue in this model.

Given this is time series data, it is likely that serial correlation is an issue. Serial correlation occurs when the errors associated with a given time period carry over into future time periods. The Gauss-Markov theorem requires simultaneous homoskedasticity and serially uncorrelated errors (Wooldridge, 2009). OLS will no longer be the best linear unbiased and efficient estimator in the presence of serial correlation. While serial correlation does not affect the unbiasedness or consistency of the OLS estimators, it does influence the efficiency. Positive serial correlation will cause the OLS estimates of the standard errors to be smaller than the true standard errors. In this case the standard errors and test statistics are not valid, even asymptotically (Wooldridge, 2009).

Given the size of the sample, the presence of serial correlation cannot be ignored. To identify serial correlation, a Durbin-Watson d-statistic is calculated. A value of 1.521 is returned. Calculating the upper and lower bounds, 0.998 and 1.931 are obtained respectively. As the DW statistic lies between the upper and lower bounds, the test is inconclusive. A DW value close to 2 suggests that autocorrelation may not be a problem.

Given a value of 1.5 and an inconclusive test, the data will be treated as if serial correlation is present, although formally the null hypothesis cannot be rejected.

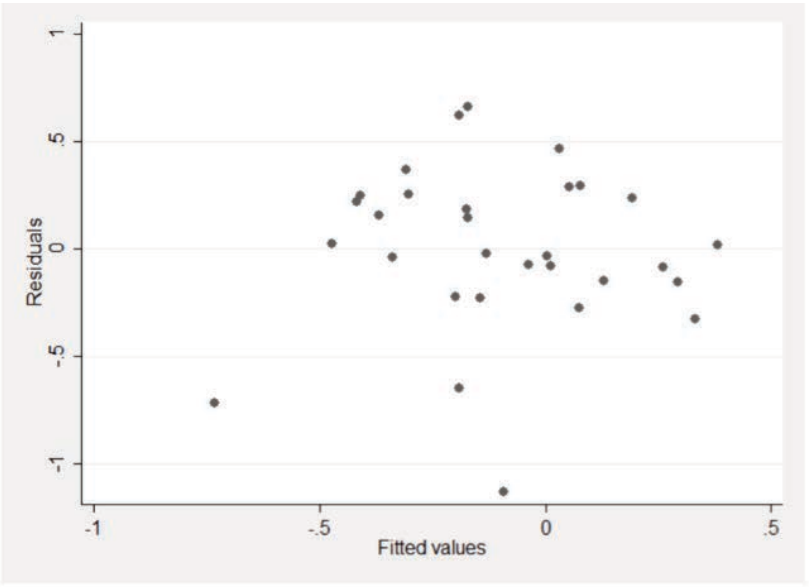

Figure 2: Residual and Fixed Value Plot

A common assumption underlying time series analysis is that the data is stationary. Many economic time series, however, are not stationary and may be highly persistent. There is reason to believe that the variables exhibit a unit root in their time series representations. A Dickey-Fuller test for unit root is conducted on the dependent and independent variables. The null hypothesis is that the variable follows a unit-root process and is therefore non-stationary. The Dickey-Fuller tests reveal that both the dependent and independent variables all exhibit unit root processes and are non-stationary.

To correct for the unit root processes, the variables are first-differenced. First differencing involves transforming the variables into a series of changes from one period to the next. First differencing can also be used to correct for the serial correlation identified earlier. The differenced data will contain one less data point than the original sample.

The regression is estimated again using the first-differenced variables as in Equation 2. Even though several variables were found to be insignificant in the first regression, they are important control variables and are retained in the revised model. Table 3 outlines the results from the regression using the first-differenced variables and robust standard

errors. A Durbin-Watson d-statistic of 2.28 is calculated. Given this is close to 2 we can conclude that the serial correlation issue has been addressed.

$$Y_i = \beta_0 + \beta_1 \Delta X_1 + \beta_2 \Delta X_2 + \beta_3 \Delta X_3 + \beta_4 \Delta X_4 + \beta_5 \Delta X_5 + \beta_6 \Delta X_6 \qquad (2)$$

From Table 3, the differences in unemployment, VAT, excise, and 25-35 population proportion are now all insignificant. The difference in the employment to population ratio is significant at the 10 per cent significance level, while the difference in log GDP per capita is significant at 5 per cent. The R-squared shows that the model is now explaining about 31.9 per cent of the difference in beer consumption, down from 92.36 per cent in the previous regression.

A negative coefficient on the employment to population ratio variable supports our hypothesis. A 1 per cent increase in the employment to population ratio will lead to an estimated decrease in litres of alcohol consumed, corresponding to beer, by about 0.43 per cent. That implies that a falling employment to population ratio, as would be expected in a recession, would lead to an increase in the consumption of beer.

A positive coefficient on the GDP variable is as expected and suggests a pro-cyclical relationship between beer and the business cycle validating the income effect. An increase of 1 in the log of GDP/capita will increase litres of alcohol in beer consumed by 8.178. Given this variable was transformed by a logarithm, we can interpret this as a percentage increase. A 1 per cent increase in GDP/capita is estimated to increase litres of alcohol consumer, corresponding to beer, by about 0.081 per cent.

Beer	Coefficient	Robust Std. Err.	t	p-value
Unemployment	-0.237	0.211	-1.13	0.272
Log GDP per capita	8.178	3.846	2.13	0.044
Employment to Population	-0.423	0.235	-1.8	0.085
Excise	0.745	3.056	0.24	0.810
VAT	2.469	4.031	0.61	0.546
25-34 population proportion	27.949	40.740	0.69	0.5
Constant	-0.362	0.178	-2.03	0.054

Table 3: OLS Regression with First Differenced Variables

Extensions

The main limitation in this study is the size of the sample. The 31 annual data points, representing 1983-2013, may not be sufficient to identify changes in demand for beer to changes in economic performance variables. Increasing the frequency of the data to allow within-year variations in alcohol consumption to be identified would be an immediate improvement. However, there are limitations to the data available and aggregated data is often the only option.

Regarding the demographic factors affecting beer consumption, the distribution of ages in the population may affect beer consumption in more nuanced ways than the single age group variable included in this study can identify. Even though we found the proportion of individuals aged 25-34 to be an insignificant variable, there may be merit in examining changes in the entire age distribution of the population. Freeman (2011) takes this approach and uses a polynomial distributed lag (PDL) model as an alternative to the single age cohort variable.

Furthermore, it would be interesting to investigate how changing economic conditions impacts the alcohol consumption of different subgroups of the population. Details on income and education could provide insights here. If lower income groups are more likely to face job losses, they may increase consumption. However, high income earners may be more exposed to other factors affected by detrimental economic performance, consider investments in property or equities. The role of gender could be examined. Bor *et al.* (2013) found single men are most likely to increase drink intake during a recession.

Finally, some literature claims that the consumption of beer rises in a recession, as consumers switch away from more expensive drinks. Further research to identify trends across alcoholic beverages could be undertaken.

Conclusion

This paper attempted to identify the cyclicality of Irish beer consumption, in particular, the hypothesis that the consumption of beer is acyclical was tested. This research was motivated by psychological work which has shown that individuals may self-medicate with alcohol and rely on it as a 'cold comfort' to alleviate the stresses of an economic downturn and associated insecurities regarding income and employment. While research has been conducted on this topic, no studies have examined the Irish experience, especially relevant giving varied business cycles across the 1980s, 1990s, and 2000s, and one of the highest beer consumption per capita statistics in the world-98.3 litres in 2012.

The findings in this paper do not lend themselves to any firm conclusions. While a negative relationship was established between beer consumption and the employment to population ratio, a positive relationship was also recorded with GDP per capita. Overall,

and following the results of the majority of the literature, it is likely that beer is a normal pro-cyclical good and that the income effect dominates any provocation effect which may increase consumption in certain circumstances.

References

Blake, D. and Nied, A. 1997. 'The demand for alcohol in the United Kingdom'. Applied Economics, 29:12:1655-72.

Bor, J., Basu, S., Coutts, A., McKee, M. and Stuckler, D. 2013. 'Alcohol Use During the Great Recession of 2008-2009'. Alcohol and Alcoholism, 48:3:343-48.

Brenner, M. and Mooney, A. 1983. 'Unemployment and health in the context of economic change'. Social Science & Medicine, 17:16:1125-38.

Dee, T. 2001. 'Alcohol abuse and economic conditions: Evidence from Re- peated Cross-sections of Individual-level Data'. Health Econ, 10:3:257-70.

Freeman, D. 2001. 'Beer and the Business cycle'. Applied Economics Letters, 8:1:51-54.

Freeman, D. 2011. 'Beer in Good Times and Bad: A U.S. State-Level Analysis of Economic Conditions and Alcohol Consumption'. Journal of Wine Economics, 6:2:231-51.

Kerr, W., Greenfield, T., Bond, J., Ye, Y. and Rehm, J. 2004. 'Age, period and cohort influences on beer, wine and spirits consumption trends in the US National Alcohol Surveys'. Addiction, 99:9:1111-20.

Tremblay, V. and Tremblay, C. 2005. The U.S. brewing industry. Cambridge, Mass: MIT Press.

Wooldridge, J. 2009. Introductory econometrics. Mason, OH: South Western, Cengage Learning.