A Bayesian Look at Consumption

PATRICK HONOHAN*
Central Bank of Ireland

Précis: This paper assesses various possible explanations for savings behaviour in Ireland during the 1960s and 1970s. Attempts to discriminate between these rival hypotheses have been hampered by multicollinearity. This problem is tackled here by the addition of prior information along Bayesian lines, using the "contract curve" framework recently suggested by Leamer.

I INTRODUCTION

It is time to have another look at the consumption function for Ireland. MODEL-80, the latest version of the Central Bank/Department of Finance macroeconometric model (Bradley *et al.*, 1981), has gone back to a simple consumption function without any of the exotic features suggested over the years. This function does not fit particularly well, but neither do the more complicated ones which have been estimated.

There are two new sources of impetus in this field. One is the newly available quarterly national accounts estimates prepared by O'Reilly (1981). In the present paper we do no use these data: another paper will do so. Here we concentrate on methodology and confine ourselves to the existing annual data. The methodological innovation is the use of the Bayesian approach suggested by Leamer (1978) and implemented in the computer program SEARCH (Leamer (1977)).

This is not the place to go into a full description of Leamer's approach. This may be found in the above references or, more succinctly, in Leamer

and Leonard (1979). It is hoped that enough will be said in what follows to allow the reader to obtain a sufficient picture of what is going on and allow the results to be readily interpreted.

II A SYSTEMATIC APPROACH TO CONSUMPTION FUNCTION ECONOMETRICS

In this study of Irish consumption data, our intention is to throw some light on recent controversies concerning the causes of wide swings in the savings ratio in the 1970s (Figure 1). The literature advances several independent but not mutually exclusive explanations for these swings. From an econometric point of view, each of these explanations amounts to a hypothesis involving an additional explanatory variable in the equation determining the savings ratio. The problem up to now has been that these rival explanatory variables have been quite highly correlated, making it difficult to choose between them. As always, in the case of highly correlated explanatory variables, the only way to achieve well defined parameter estimates is to take account of prior information, either explicitly, as we will do here, or implicitly, as is done in, for example, ridge regression techniques.

One way of looking at this problem is to consider it from a classical point of view in which the model includes all the explanatory variables which have been proposed, but from which different estimates have been obtained by different researchers imposing different prior exclusion restrictions. Now the set of all the estimates that could be obtained subject to any generalised linear exclusion restrictions (i.e., restrictions on the parameter vector β of the form $R\beta=0$) is a well defined ellipsoid in parameter space. No matter what exclusion restrictions a researcher places on the parameters, his (least squares) estimates must be within this ellipsoid, referred to by Leamer as the *"set of constrained least squares points"*. This set is obviously of interest to the reader of econometric results who does not necessarily share the researcher's views as to the appropriate parameter restrictions.

A Bayesian interpretation of this "set of constrained least squares points" finds it to be equivalent to the set of all possible modes that might be generated by different researchers with different prior distributions so long as those priors were multivariate normal with mean zero. Thus, so long as the mean of the prior distribution is at the origin (and the prior distribution is multivariate normal[1]) the mode of the posterior distribution will be within the set of constrained least squares points. This set can be calculated from knowledge of the data alone, and is independent of how strongly or weakly

---

1. In fact, the prior distribution need only be elliptically uniform, and even this requirement can be weakened.

one feels about one's prior beliefs. It is also independent of one's prior correlations between the values of different parameters.

Figure 1: *Savings Ratios 1960-79*
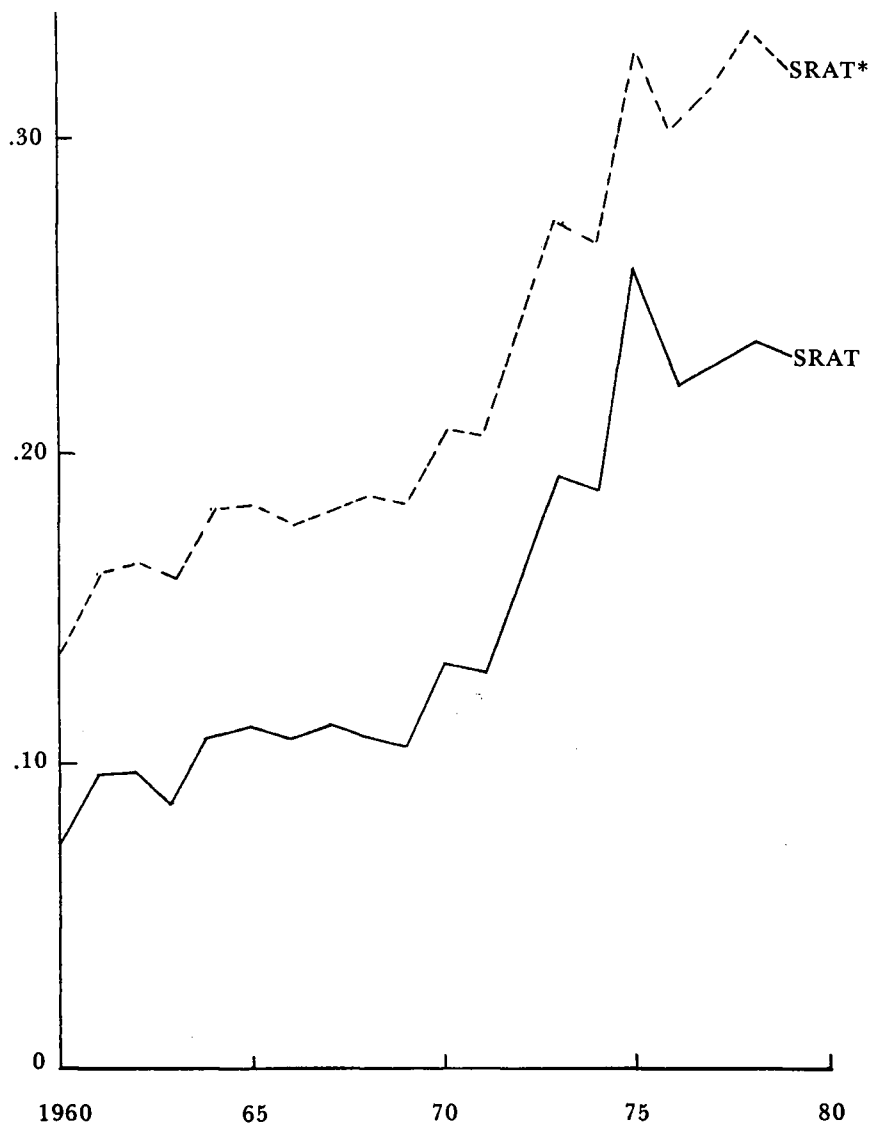
SRAT — Personal savings as proportion of personal disposable income.
SRAT* — Personal savings plus expenditure on consumer durables as proportion of personal disposable income.

In case the prior distributions have mean other than zero, we can interpret the above comments as referring to the transformed parameter set $y = \beta - \beta^*$ where β^* are the prior means.

Although the set of contrained least squares points is thus of great interest in principle, we will find little use for it in the sequel, since we will discover that it is a large set given the data we are using. This means that restricting ourselves to parameter values within the set does not in fact constrain the parameter values very much. The set does not greatly narrow the scope of our attention.

In order to focus on a subset of the set of constrained least squares points, we must, therefore, be prepared to commit ourselves to more prior information than simply the information that our prior is multivariate normal with a specified mean vector β^*. Nevertheless, we do not have to specify the prior completely. We can leave a scalar confidence factor unspecified and still narrow our focus very considerably. The result is a curved line in the parameter space known as the "information contract curve" (Figure 2). This is the set of all points which could be modes of the posterior distribution given the data and the prior means and covariance matrix for different researchers whose prior covariance matrices were scalar multiples of each other. Thus, if we can specify the structure of prior information then, even without knowing the degree of confidence to be placed in prior information by comparison with the data information, we can deduce that the posterior mode will be on the information contract curve.

Our results will focus on the information contract curve — or simply the "contract curve" for short. In addition to presenting plots of the contract curve, we will attempt, in an informal fashion, to deduce *ex post* a plausible range for the prior confidence factor — the scalar which pinpoints the exact location of the posterior mode on the contract curve.

III THE EMPIRICAL ISSUES

After hovering around 11 per cent for most of the 1960s, the savings ratio in Ireland jumped to 13 per cent in 1970, to 16 per cent in 1972, to 19 per cent in 1973 and by an extraordinary seven points to almost 26 per cent in 1975. From 1976 to 1979 the ratio remained between about 22 and 23½ per cent. Our data ends in 1979, but preliminary estimates for 1980 and 1981 suggest a further fall in the savings ratio in these years to below 20 per cent.

Why did the savings ratio rise so sharply in the mid-1970s and, in particular, how are we to interpret the outturn for 1975? In that year inflation reached a record level, as did the unemployment rate, which rose by one-third. The real value of liquid assets had never been lower in proportion to disposable income. Each of these developments has been associated — by

different authors — with the sharp rise in the savings ratio in 1975. Analysis of the data reveals that any one of these variables — inflation, unemployment, liquidity ratio — can indeed explain a good deal of the variation in the savings ratio over the past twenty years[2]. The problem is that the high correlation between these variables introduces severe multicollinearity, making it difficult to distinguish between them and assign relative degrees of importance.

Figure 2: *The Information Contract Curve*
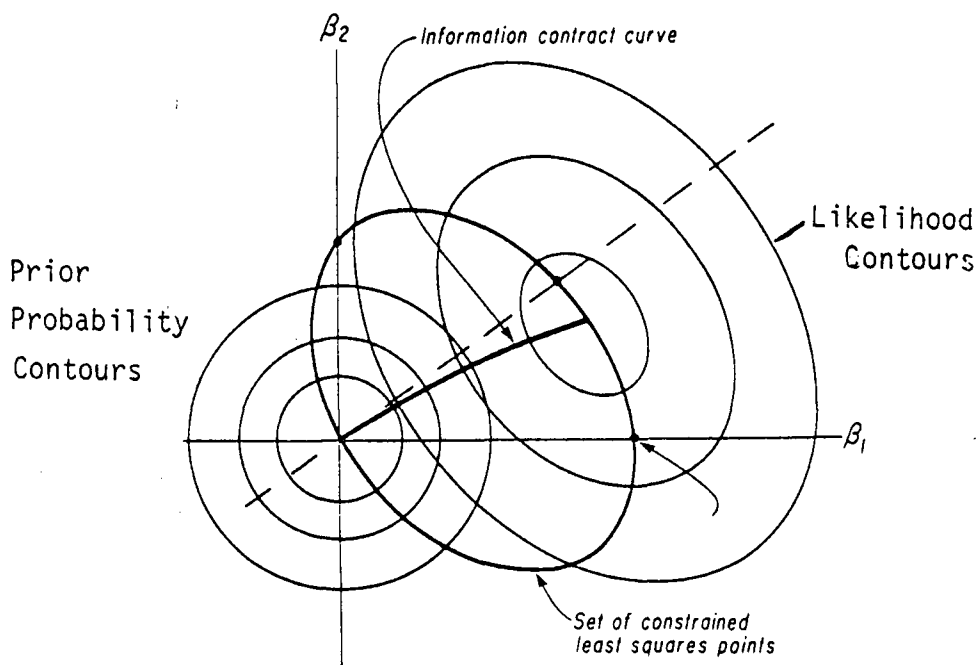
Multicollinearity and a shortage of degrees of freedom also creates other difficulties. Of particular importance in this context is the estimated short-run marginal propensity to save (SRMPS). This coefficient is often not well determined and point estimates are often implausible. Furthermore, it is difficult to assess the importance of other variables which have been thought relevant to the consumption function.

The addition of extra information is the only real solution to problems caused by multicollinearity. In the absence of additional data, prior informa-

2. Digby (1980) is a good review of the various models which have been proposed and their relative explanatory power. More recently, unpublished studies by G. Boyle and by M. O'Mahony have examined a further theory which asserts that capital losses on private sector financial assets should be netted out of income for the purposes of estimating the consumption function. This theory too relies on inflation and the liquidity ratio as explanatory variables, since their product is the capital loss on financial assets relative to income.

tion consistently applied according to Bayesian methodology may be used to clarify the issues raised here. The specification of this prior information is the subject of the next section.

IV VARIABLES AND PRIORS

Our approach in specifying the model has been to draw on previous research for a list of explanatory variables which have a good claim to inclusion in the consumption function. We have not devoted a great deal of attention to the choice of functional form, as this is not easily handled with the methodology and software being employed here. Nor have we elaborated the dynamics to the extent that would be required, for example, in a quarterly study.

A reading of the literature suggests nine explanatory variables which should be examined. These may be grouped into four *categories* as follows:
  (i)   functional form;   F
  (ii) neo-classical effects; N
  (iii) disequilibrium;   D   .
  (iv) surprise.   S

  (i) The first two variables may be grouped together as representing modifications to the functional form rather than rival economic theories. It has often been argued (e.g., by Kennedy and Dowling (1970)) that consumption out of agricultural incomes is less than consumption out of non-agricultural incomes. The rationale could be that agricultural incomes are more variable, so that Friedman's permanent income hypothesis would lead to a greater errors-in-variables problem with agricultural consumption. In any event, an intercept shift proportionate to the ratio of agricultural income to total personal disposable income has been used to capture such a differential effect in the past, and the same is done here. This variable is denoted AGRINCOM.[3] The other functional form variable is real personal disposable income. Since our dependent variable is the savings ratio, this represents a quadratic term in income for consumption as a function of income; we call it SQINCOM.
  (ii) The second group of variables are described as neo-classical variables, in that they are prices to which even an unconstrained agent would be expected to react. One (found in Bradley (1979)) is the ratio of non-durable to overall consumer prices (RELPRICE) — relevant to the choice between non-durables and durables — and the other is the real

3. Variables are defined in the data appendix.

interest rate (REALINT) — relevant to the choice between consumption now and consumption later.

(iii) The third category of variables captures various sources of slow adjustment in the economy. A liquidity variable (LIQUID) should capture the reduction (increase) in saving which will follow an unanticipated change in the real value of liquid assets (Kelleher (1977)). (Actually, we normalise liquid assets by real disposable income; the target for this ratio seems more likely to be stable than for real liquid assets as such). If liquidity was always on target, there would not be this gradual stock adjustment which has consequences for saving. The unemployment rate (UNEMPLOY) is included to take account of the influence of slow-clearing labour market conditions on saving (McCarthy (1979)). The effect of unemployment on saving is not entirely clear a priori. On the one hand, it may reduce the savings ratio since it means that more of personal disposable income is represented by transfers; on the other hand, a high unemployment rate may increase the savings ratio as it induces precautionary saving on the part of those who, though not unemployed, perceive the risk of unemployment as being higher and thus their (expected) permanent income as being lower. The lagged savings ratio (LAGSR) directly measures slow adjustment in consumption, though it could also be implicitly making an imperfect correction for auto-correlation in the residuals.

(iv) The final category of explanatory variable represents surprises or rationing. Both variables, the rate of inflation (INFLAT) and the real rate of growth of personal disposable income (GROWTH) derive from Deaton's (1977) hypothesis (Honohan (1979)). It is possible, however, that the inflation rate explains variations in consumption behaviour for other reasons, such as those mentioned in footnote 2.

V CONSTRUCTING THE PRIOR

The question of how to construct a prior distribution which would be believable, and so allow the empirical analysis to be of general interest, raised difficult conceptual and practical puzzles. Our approach was to try to approximate, in the prior, the fuzzy notions held by the average reader of the literature on consumption behaviour in Ireland. We assumed that such a reader is inclined to place a non-zero degree of trust in the methodology and skill of the various researchers. The reader combines the results in an informal way and ends up with beliefs about the consumption function which are a loose amalgam of the researchers' findings.

It must be recognised that such beliefs could hardly be justified from a formal statistical analysis of the data. Indeed, if they could be, it would be

difficult to draw inferences from the present exercise which is, after all, based on much the same data[4] as used by previous researchers. Instead, the beliefs are heavily influenced by the way in which earlier studies have been reported and thus will typically depend on the researchers' models, implicit prior beliefs and methodologies.

Accordingly, we based the prior *means* and relative *standard errors* on the reported findings of previous researchers (as cited above, see also Bradley *et al.*, (1981) and Digby (1980)). While the exact choices listed in Table 1 are necessarily somewhat subjective, it may be remarked that the researchers concerned did not raise objections to these choices as representing their findings when they commented on earlier drafts of this paper! To assist in the assessment of these prior means and standard errors, which are specified for a linear model, approximate elasticities calculated at data means are also shown in Table 1.

Table 1: *Prior means and (relative) standard errors*

Category	Variable	Prior mean	Standard error	Prior mean elasticity*	Standard error*
F Functional Form	SRMPS	0.50	0.50		
	AGRINCOM	0.15	0.15	0.20	0.20
	SQINCOM	0.00	1.00	0.00	1.60
N Neo-classical	RELPRICE	0.00	0.40	0.00	2.70
	REALINT	0.00	0.50	0.00	0.30
D Dis-equilibrium	LIQUID	-0.75	0.15	3.50	0.70
	UNEMPLOY	1.30	0.30	0.40	0.10
	LAGSR	0.50	0.10	0.50	0.10
S Surprise	INFLAT	0.40	0.10	0.40	0.10
	GROWTH	0.40	0.10	0.13	0.03

*Approximate: computed at data means, dependent variable is saving ratio SRAT.

4. Our data are, however, more up to date than in most of the studies cited.

In order to construct the prior *correlation matrix* we relied on rather different considerations than for the mean and standard error. Recalling the classification of explanatory variables into four categories, we first considered prior correlations between categories (cf. Table 2; each category is labelled by a letter F,N,D,S).

Table 2: *Breakdown of prior correlation matrix*

(i) *Within categories*

F | 1
   | 0   1

D | 1
   | .4   .1
   | .2   -.2   1

N | 1
   | .2   1

S | 1
   | .4   1

(ii) *Between categories*

	F	N	D	S
F	*			
N	0	*		
D	0	-.1	*	
S	0	-.2	-.4	*

(iii) *With SRMPS*

D: -.1, -.1, -.2;   S: -.2, -.2

Note: The letters F, N, D, S refer to the categories of explanatory variable described in text, Section IV.

There seemed no reason to assume a non-zero correlation between non-linearity coefficients and the rest. A weak negative correlation between the neo-classical coefficients and the remainder was postulated on the ground that the greater the effect of neo-classical forces, the less one would look to disequilibrium or surprise variables for an explanation of fluctuations in saving. A relatively high negative correlation between the disequilibrium coefficients and the surprise coefficients reflects the fact that these two classes of theory have been advanced as rival explanations for the savings behaviour of 1974-1976. So far as the column representing the SRMPS is concerned, we proposed negative correlation between this and some of the other coefficients, since these variables have been introduced by research-ers to explain exceptionally high observed savings ratios. It is reasonable to suppose *ex ante* that, the stronger these effects, the lower the actual SRMPS.

Having settled the prior correlations between categories, it remained to fix the correlation within categories. These sub-matrices are shown in Table 2.

Since the two non-linearity variables are unrelated, a zero prior correlation for their coefficients was proposed. If neo-classical factors are important for consumption, it is plausible to suppose *ex ante* that both intertemporal sub-stitution and substitution between classes of commodities could be relevant. Knowledge that one type of substitution is important would increase one's expectation of the size of the other substitution. A moderate positive corre-lation between these two was entered.

In much the same way we took it that the magnitudes of the disequili-brium factors were correlated *ex ante*. The sign of this correlation was determined by the prior mean as to the effects of these variables. Further-more, we imposed a higher correlation between the indirect disequilibrium factors LIQUID and UNEMPLOY than between either of these and the (direct) lagged savings ratio. The relatively high correlation imposed between the two "surprise" terms reflects their common origin in Deaton's theory.

Choice of prior distributions is always a tricky problem for Bayesian analysis. If the researcher cannot convince the reader of his choice then the reader will not be prepared to accept the conclusions. However, at least in using the contract curve the reader need not accept the researcher's degree of confidence. The prior outlined in Tables 1 and 2 was not chosen by working back from some prejudices about the results. Indeed the complexity of attempting such a construction would seem likely to deter any unscrupulous researcher from doing so. It is hard enough to ensure that one's prior covar-iance matrix is positive definite!

A reasonably modest way of interpreting the contract curve based on this prior (and on the data) is to say that it tells us the direction in which the data can be most easily coaxed if faced with some more or less vague prior information. All the functional relationships involved are continuous ones

and a small change in the prior parameters will not mean a large jump in the contract curve. In any event, as already mentioned, so far as annual data are concerned, adding prior information is the only way forward out of the present difficulty of multicollinearity. The fairly systematic manner in which the prior has been specified should enhance our confidence in the relevance of the estimated contract curves.

VI THE RESULTS

We concentrate on results for the dependent variable SRAT* which is the savings ratio where expenditure on consumer durables is included in savings[5].

The contract curve is reported by plotting its projection onto each parameter co-ordinate axis in Figure 3. On the right-hand edge of each panel in Figure 3, we see the "data point", i.e., the unconstrained LS coefficient estimate based on a linear regression of SRAT* on all the explanatory variables using annual data 1961-1979. For instance, the unconstrained estimate of the relative price coefficient RELPRICE is 0.5. As we move left in each panel, we trace out the posterior modes we would obtain if we had more confidence in our prior relative to the data. The base scale shows the confidence intervals which the contract curve intersects at each point. For example, estimates to the left of 0.95 are outside the 95 per cent confidence ellipsoid and would (in the absence of prior information) be rejected at the 5 per cent level of significance; while points just to the left of 0.85 could only be rejected at the 15 per cent level.

Readers can choose their preferred level of significance, or equivalently the scale factor by which they would wish to multiply all the prior (relative) standard errors shown in Table 1. Table 3 compares the LS data point with the prior, using a scale factor of 2, and the resulting posterior mode (again using scale factor 2). This posterior is just outside the 85 per cent confidence ellipsoid[6].

We may use the contract curves, and posterior standard errors, to interpret our findings as follows. For coefficients where the contract curve slopes steeply on the left-hand side, near the prior mean, but is flat near the LS data point, the prior information conflicts sharply with the data, in respect of the coefficient. Only a high degree of confidence in the prior information would allow us to modify the LS data point. On the other hand, for coefficients where the contract curve is flat near the prior but slopes steeply

5. The results for the more conventional savings ratio SRAT not including expenditure on durables are quite similar; details are available from the author.
6. For those who are not happy with the prior choice of means, the derivative of the posterior (scale factor 2) means with respect to prior means for the significant coefficients are: SRMPS .14; LIQUID .16; UNEMPLOY .23; LAGSR .56; INFLAT .17; GROWTH .27.

near the data point, we can conclude that multicollinearity is a problem and that the data do not need much "persuasion" by prior information to have the posterior much closer to the prior mean.

Figure 3: *The Contract Curve*
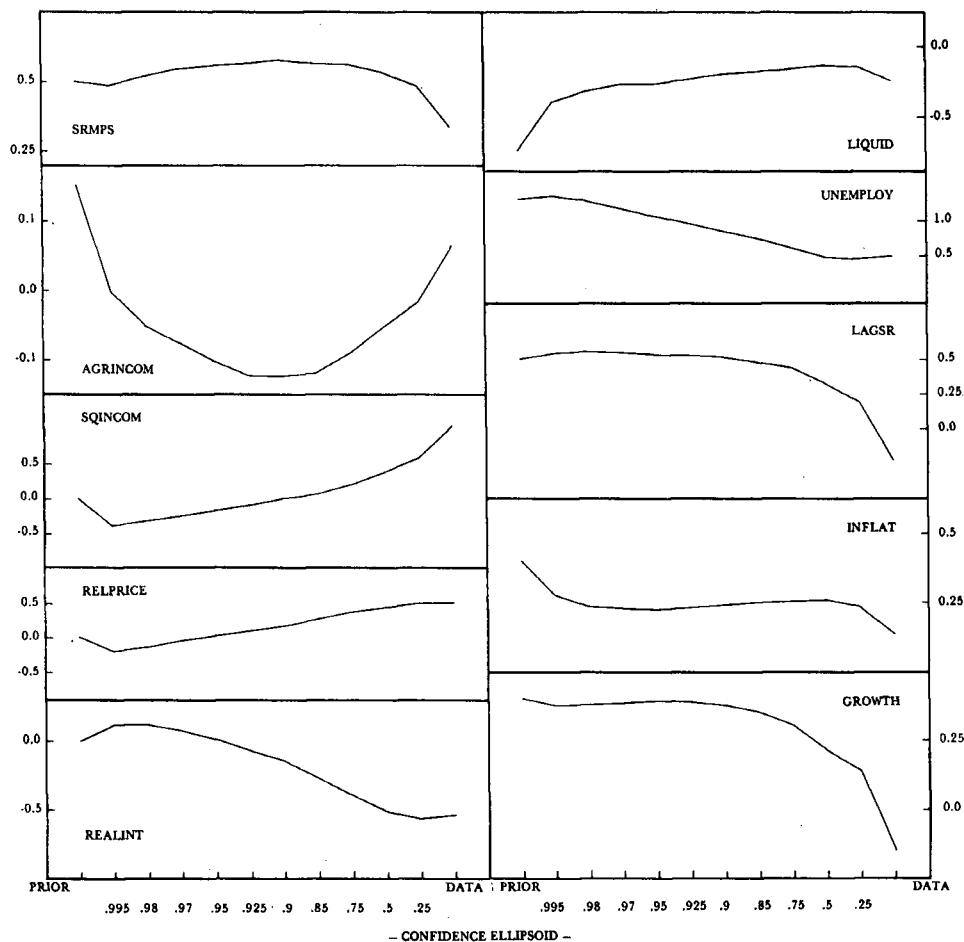
The coefficients for LIQUID, and to a lesser extent UNEMPLOY fall into the first type; SRMPS, LAGSR and GROWTH into the second type; while INFLAT has features of both types. We conclude that multicollinearity is biasing the LS estimates of the coefficients for SRMPS, LAGSR, GROWTH and INFLAT —

the short-run marginal propensity to save is much closer to its prior value of 0.5 than to the LS value of 0.3;

the adjustment coefficient on the lagged dependent variable LAGSR is

positive and close to its prior value of 0.5, rather than the negative LS value;

likewise the GROWTH coefficient is positive rather than negative;

the best estimate of the INFLAT coefficient is much closer to 0.25 than to either the LS estimate of 0.15 or the prior mean 0.4.

Table 3: *Estimated coefficients (with standard errors in parentheses)*

Dependent variable: SRAT*	LS estimates	Prior (scale factor 2)	Posterior
SRMPS	0.332* (0.084)	0.50* (1.00)	0.577* (0.049)
AGRINCOM	0.064 (0.205)	0.15 (0.30)	-0.122 (0.158)
SQINCOM	1.063* (0.263)	0.00 (2.00)	0.055 (0.161)
RELPRICE	0.478* (0.223)	0.00 (0.80)	0.240 (0.195)
REALINT	-0.543 (0.355)	0.00 (1.00)	-0.265 (0.283)
LIQUID	-0.243 (0.130)	-0.75* (0.30)	-0.205* (0.092)
UNEMPLOY	0.486 (0.343)	1.30* (0.60)	0.742* (0.264)
LAGSR	-0.233 (0.203)	0.50* (0.20)	0.496* (0.094)
INFLAT	0.146* (0.083)	0.40* (0.20)	0.248* (0.057)
GROWTH	-0.148 (0.122)	0.40* (0.20)	0.357* (0.052)
$\bar{R}^2$	0.991		
DW	1.55**		
Sample	Annual 1961-1979		

*Note: Coefficient values greater than twice the corresponding standard error are marked with an asterisk.

** Serial independence not rejected by Durbin's residual regression test.

On the other hand, the degree of prior confidence required to move the posterior modes for LIQUID and UNEMPLOY far from their LS values suggests that for these coefficients the best estimate is closer to the LS value than to the prior.

The addition of prior information also reduces the posterior standard errors. This does not, however, increase the posterior significance of all variables. As can be seen from Table 3, the addition of prior information enhances the significance of SRMPS, LIQUID, UNEMPLOY, LAGSR, INFLAT and GROWTH and we may conclude that these variables are all significant determinants of savings behaviour. But prior information diminishes the significance of the other variables. Except for AGRINCOM, this is because the prior mean is zero, but it does seem that these variables may not have a significant effect on savings behaviour despite the apparent significance of two of them in the LS estimates.

It may reasonably be asked: what consumption function should be used as a result of this analysis? One answer to this question might be: the best function for each reader is that given by the contract curve at the point corresponding to the reader's preferred prior confidence scale factor. Now it may be difficult to decide *ex ante* on what scale factor corresponds to one's prior beliefs. If so, the following reasoning may assist in suggesting an appropriate scale factor *ex post*. Our purpose in adding prior information was to resolve a deficiency of data manifested in multicollinearity. Our prior, if adequate, should have resolved this multicollinearity by causing the posterior to be significantly different from the data point, but without being dominated by the prior information. Consulting the contract curve in Figure 3 we would deduce that, were our posterior at 0.25, then the prior information would have been insufficiently precise to resolve the multicollinearity problem, since the contract curve is quite close to the data point at 0.25. On the other hand, a posterior to the left of 0.995 would be ignoring the data to an extent clearly unwarranted by the degree of confidence which one could reasonably place in the prior construction. The posterior reported in Table 3, corresponding to 0.85 on Figure 3 might seem a reasonable compromise in this light.

VII CONCLUSION

A slavish acceptance of posterior modes is not, however, essential. Guided by the posterior standard errors we may draw the following general conclusions:

    none of the explanatory variables specifically introduced to explain the unusual savings behaviour of the 1970s (inflation, the real level of liquid

assets, real income growth and the rate of unemployment) can be dismissed as irrelevant though their coefficients are smaller than we may have believed on the basis of previous work;

some other variables (the share of agricultural income, the interest rate, relative prices of durables and non-durables and a quadratic term in income) do not appear to be significant influences on savings behaviour;

the short-run marginal propensity to consume seems to be about one-half.

REFERENCES

BRADLEY, J., 1979. "A Note on the Estimation of Disaggregated Consumption Functions", *Technical Paper No. 8/RT/79*, Dublin: Central Bank of Ireland.

BRADLEY, J., C. DIGBY, J.D. FITZGERALD, O. KEEGAN and K. KIRWAN, 1981. "Description, Simulation and Multiplier Analysis of the MODEL-80 Econometric Model of Ireland, *Research Paper 20/81*, Dublin: Department of Finance.

DEATON, A., 1977. "Involuntary Saving Through Unanticipated Inflation", *American Economic Review*, Vol. 67, No. 5, pp. 889-910.

DIGBY, C., 1980. "A Review of Consumption Studies in Ireland", Mimeo, Central Bank of Ireland.

HONOHAN, P., 1979. "A Note on Inflation and Consumption in Ireland", *The Economic and Social Review*, Vol. 10, No. 2, pp. 175-179.

KELLEHER, R., 1977. "The Influence of Liquid Assets and the Sectoral Distribution of Income on Aggregate Consumers' Behaviour in Ireland", *The Economic and Social Review*, Vol. 8, No. 3, pp. 189-199.

KENNEDY, K.A. and B.R. DOWLING, 1970. "The Determinants of Personal Savings in Ireland: An Econometric Inquiry", *The Economic and Social Review*, Vol. 2, No. 1, pp. 19-51.

LEAMER, E.E., 1977. "SEARCH, A Linear Regression Computer Package", *Harvard Institute for Economic Research Discussion Paper 89*, March. (Revised January 1979).

LEAMER, E.E., 1978. *Specification Searches*, New York: Wiley.

LEAMER, E.E. and H. LEONARD, 1979. "An Alternative Reporting Style for Econometric Results", *Harvard Institute for Economic Research Discussion Paper 145*, February.

McCARTHY, C., 1979. "A Study of Aggregate Consumer's Behaviour in Ireland", Mimeo, Central Bank of Ireland.

O'REILLY, L., 1981. "Estimating Quarterly National Accounts", Forthcoming in the *Journal of the Statistical and Social Inquiry Society of Ireland.*

DATA APPENDIX

The variables used in the analysis are set out below, together with their definitions. Except for M3, which is the revised wide money series published in the *Central Bank Annual Report 1982,* all variables are constructed from the Department of Finance/Central Bank Annual Databank and the definitions given below use the mnemonics of that databank.

Dependent Variables

SRAT*   Personal savings plus expenditure on consumer durables as proportion of personal disposable income.

$$1 - (CND*PCND + CS*PCS)/YDV$$

SRAT   Personal savings as proportion of personal disposable income.

$$1 - (CND*PCND + CS*PCS + CD*PCD)/YDV$$

Functional Form Variables

AGRINCOM Agricultural income as proportion of personal disposable income.

YA/YDV

SQINCOM   Real personal disposable income ($£10^{10}$).

$(YDV/PC) \times 10^{-4}$

Neo-classical Variables

RELPRICE Ratio of non-durable to overall consumer prices.

PCND/PC

REALINT   Real interest rate.

$(IPLR - PCDOT) \times 10^{-2}$

Disequilibrium Variables

LIQUID   Wide money (end of previous year) as a proportion of personal disposable income.

M3(-1)/YDV.

UNEMPLOY Unemployment rate.
          $UR \times 10^{-2}$

LAGSR   Lagged dependent variable.

Surprise Variable
INFLAT   Rate of consumer price inflation.
             PCDOT

GROWTH   Rate of growth of real personal disposable income.
             (YDV/PC)/(YDV(-1)/PC(-1))-1

It seems appropriate to note in this appendix that some economists have suggested privately that the remarkable behaviour of the savings ratio in the 1970s could be attributable to some systematic data errors in the consumption and/or income data. We have not been able to find firm evidence on this conjecture.