Mind your ps and qs! Improving ARMA forecasts with RBC priors

Kirdan Lees and Troy Matheson

Reserve Bank of New Zealand

JEL classification: C11, C22, E37

2005

Preliminary, Work in Progress

Mind your ps and qs! Improving ARMA forecasts with RBC priors*

Kirdan Lees and Troy Matheson†

Abstract

We test whether information from an RBC model can improve ARMA forecasts of post-war US GDP. We find the RBC model selects ARMA orders that yield improved forecasting performance over conventional information criteria. Using a Bayesian framework to take the model to the data we find comparable forecasting performance between the RBC model and the ARMA model estimated using maximum likelihood.

*We thank Shaun Vahey and seminar participants at the Reserve Bank of New Zealand for comments. Any remaining errors are the responsibility of the authors. The views in the paper are those of the authors and these views should not be attributed to the Reserve Bank of New Zealand.

†Correspondence Address: Troy Matheson, Economics Department, Reserve Bank of New Zealand, P.O. Box 2498, Wellington, New Zealand; Tel: +64 4 471 3859, Fax +64 4 473 1209, Email: troy.matheson@rbnz.govt.nz

1 Introduction

Our computational experiment evaluates whether a minimal set of RBC theory can improve forecasting. If we view a statistical ARMA model as potentially misspecified it is an empirical question whether information encapsulated in an RBC model is informative for forecasting purposes.

Prior information from DSGE models has been shown to be useful in improving forecast performance of VAR models. Ingram and Whiteman (1994) show that the forecasting performance of a Bayesian VAR with a Minnesota prior and a Bayesian VAR with an RBC prior is comparable; DeJong, Ingram, and Whiteman (2000) show that the forecasting performance of the RBC model of Greenwood, Hercowitz, and Huffman (1988) is comparable to that of a Bayesian VAR; Del Negro and Schorfheide (2004) find information in a new-Keynesian prior improves VAR forecasts for US output growth, inflation and the Federal Funds rate.

However, VARs restrict representations of data to auto-regressions with no moving average terms. While an infinite order VAR can approximate an ARMA model arbitrarily well, VARs typically used in applied work can fail to match the structure implied by theoretical RBC or DSGE models (see for example Chari, Kehoe, and McGrattan (2005), Kapetanios, Pagan, and Scott (2005) and Fernandez-Villaverde, Rubio-Ramirez, and Sargent (2005)). If these theoretical models are close to the true data-generating process (DGP), VAR models with lag lengths of typical order may fail to recover the DGP. Thus we test forecasting performance across ARMA models that contain moving average terms. Our forecasting metric for comparison is out-of-sample pseudo real-time forecasts for a measure of post-war US GDP.

We use the Campbell (1994) RBC model because it can be represented by ARMA processes for a core set of macroeconomic variables and it is parsimonious. If such a minimal set of theory helps ARMA forecasts, a wider set of models also may improve ARMA forecasts.[1] Section 2 discusses ARMA forecasting, outlines the Campbell (1994) model and presents the forecasting experiment. Section 3 presents results and section 4 concludes.

[1]DeJong, Ingram, and Whiteman (2000) make the same argument with regard to comparing forecasts from DSGE models to VAR models.

2 Forecasting with ARMA processes

2.1 Statistical ARMA estimation

Box and Jenkins (1976) motivate ARMA modelling to obtain optimal forecasts. The general ARMA model is:

$$y_t = \sum_{i=0}^{p} \phi_p y_{t-i} + \epsilon_t + \sum_{j=0}^{q} \theta_j \epsilon_{t-j} \tag{1}$$

where ϵ_t is an idiosyncratic disturbance, y_t is the stationary forecast variable.

Purely statistical univariate ARMA models often prove useful forecasting models for reasons of parsimony (Box and Jenkins (1976)). Indeed, frequently large-scale macroeconometric models fail to outperform univariate ARMA models with a small number of arguments (Hamilton, 1994, p. 109). The Box-Jenkins method has been interpreted as implying the following strategy:

1. Select an appropriate maximum p, q order for the ARMA (p,q) model.
2. Estimate $\phi(L)$ and $\theta(L)$ associated with each choice of p and q.
3. Select the best model (choice of p and q) based on model diagnostics.

Researchers are guided to appropriate maximum (positive integer) values of p and q by examining the behaviour of the ACF and PACF of the dataset to be modelled (see Enders (1995) and Hamilton (1994)) information criteria (such as the Akaike Information Criterion (AIC) and the Schwarz-Bayesian Criterion (SBC)) trade-off explanatory power with parsimony to help selection. But visual inspection may leave several models to consider.[2]

2.2 Why theory might help

We start from the supposition that both the statistical model and theoretical model misspecify the true DGP process — the model is only a summary of particular aspects of a complex real world.

Clements and Hendry (1998) form a taxonomy of matches between statistical model, theoretical model and data generating process. They note a prior derived from the theoretical model will always be helpful when the statistical model is misspecified and the theoretical model is correctly specified.

[2]Hamilton (1994), p 112-113, suggests log changes of quarterly US real GNP data 1947 to 1988 appear consistent with both AR(1) and MA(2) processes.

When the statistical model and theoretical model are both misspecified it is an empirical question whether the prior from theory will help forecasting.

For a particular ARMA process, Box and Jenkins (1976) use Bayes law to represent the posterior distribution of the ARMA parameter vector with the dataset, as the combination of prior information about the parameter vector and the likelihood of the parameter vector generating the dataset. That is:

$$p(\xi|z) = \frac{p(\xi)L(\xi|z)}{\int p(\xi)L(\xi|z)d\xi} \tag{2}$$

where ξ, is the ARMA parameter vector, z is the dataset and $L(\xi|z)$, the likelihood of the prior, $p(\xi)$, generating dataset z, such that $p(\xi|z)$ is the posterior distribution of the ARMA parameter vector.

Box and Jenkins (1976) motivate the use of Jeffrey's prior, which maintains a one to one mapping between the prior and the posterior. Jeffrey's prior is thus locally uniform and hence uninformative about the posterior. But if we acknowledge that the statistical ARMA process may be misspecified information derived from theory may improve ARMA forecasts. Bayes law gives us a consistent framework for incorporating prior information with the data for forecasting.

Furthermore, if we believe that the data-generating process is constant, then the presence of sampling error in estimates of autocorrelation coefficients suggests misspecification in the ARMA process. The choice of the order of the ARMA process varies over time. If we acknowledge the possibility that the statistical ARMA model is potentially misspecified but are willing to place some belief in theory, then this prior information can be incorporated in both the choice of the order of the ARMA process and the estimation of the ARMA process. Thus theory may prove helpful with:

1. Selecting an appropriate order for the autoregressive and moving average components (the ps and qs); and
2. Informing $\phi(L)$ and $\theta(L)$ associated with the selected model.

To test whether RBC theory is helpful, we compare forecasts from a purely statistical baseline ARMA model with three ARMA processes that incrementally appropriates additional information from the RBC model:

1. The RBC model picks the ARMA order (ps and qs); the ϕs and θs are estimated using maximum likelihood;

2. The RBC model picks the ARMA order; the RBC prior about the deep parameters of the model is combined with the likelihood using the Metropolis-Hastings simulator to generate ϕs and θs;

3. The RBC model picks the ARMA order; point mass priors for the deep parameters of the RBC model are used to calibrate ϕs and θs .

2.3 The Campbell model

Campbell (1994) appeals as an RBC model to test the ability of economic theory to improve forecasting because it is parsimonious and well understood.[3] If such a simple model helps forecasting, richer models may also improve ARMA forecasting.

We use Campbell's (1994) flexible labour model. The representative agent has log utility for consumption and power utility for leisure:

$$E_t \sum_{i=0}^{\infty} \beta^i [\log(C_{t+i}) + \theta \frac{(1 - N_{t+i})^{1-\gamma}}{1 - \gamma}] \tag{3}$$

where γ is the coefficient of relative risk aversion, and the intertemporal elasticity of substitution is defined to be $\sigma \equiv 1/\gamma$.

Firms produce with a constant returns-to-scale Cobb-Douglas function:

$$Y_t = (A_t N_t)^{\alpha} K_t^{1-\alpha} \tag{4}$$

where N_t is the number of hours of labour, K_t the stock of capital and A_t technology. Capital accumulates according to:

$$K_{t+1} = (1 - \delta) K_t + Y_t - C_t \tag{5}$$

where δ is the depreciation rate.

Campbell (1994) derives a log-linear representation of the model with an autoregressive process for the natural logarithm of technology shocks, a_t:

$$a_t = \phi a_{t-1} + \epsilon_t \tag{6}$$

[3]For example, Campbell (1994) uses the model to show that output may increase initially in the face of highly persistent negative technology shocks; Ludvigson (1996) shows that the persistence of the government debt process may obviate any crowding out effects from deficit-financed cuts in distortionary income taxation; Lettau (2003) uses a model variant to show why the equity premium is small and Campbell and Ludvigson (2001) reconcile the small intertemporal elasticity of substitution implied by the model by introducing a household sector.

Campbell shows that an analytical solution exists for the model that reduces the natural logarithm of output, y_t, to an ARMA(2,1) process:

$$y_t = (\phi + \eta_{kk})y_{t-1} - \eta_{kk}\phi y_{t-2} + \eta_{ya}\varepsilon_t - (\eta_{yk}\eta_{ka} - \eta_{ya}\eta_{kk})\varepsilon_{t-1} \tag{7}$$

where ε_t is the idiosyncratic shock to (log) technology and η_{ya}, η_{yk}, η_{ka} and η_{kk} are elasticities that are functions of the deep parameters model. The Appendix shows derivation of the ARMA representation.

2.4 The steady-state

Campbell (1994) derives the dynamic properties of the model using log-linearisations and first order Taylor approximations around steady state values. This means that the data must be expressed in terms of deviations from the steady state — but what are the appropriate steady state values?

Quantitative properties of the data can be dependant on the detrending method. Canova (1998) suggests that at a minimum, the data generated by the theory should be passed through a variety of different detrending filters to check the implications of the theoretical model over a range of frequencies. We detrend our data using two methods — we extract a linear trend (LT), and we take first order differences (FD).

3 Results

3.1 Data

We use quarterly US GDP data (in chain-linked 2000 dollars) from 1947q1 to 2004q3, obtained from the Bureau of Economic Analysis (BEA).

The calibration literature (see Cooley and Prescott (1995) and King and Rebelo (2000), for example) emphasizes matching the concept variable in the model to the corresponding concept of the variable in the data. Campbell's (1994) theoretical economy is closed, has no government sector and consumption goods only last for one period. This suggests removing net exports, removing the government sector and defining consumption narrowly, to exclude durable goods, which can exist for more than one period.

Thus we define output narrowly as the sum of private consumption and private investment, where private consumption is private non-durables consumption (PCNDGC96) plus private services consumption (PCESVC96) and private investment is defined to be total private fixed investment (FPIC1) plus private durables consumption (PCDGCC96). We also work with a broad definition of output that is total real GDP (GDPC1).

Figure 1: US Output Definitions - 1947:1 – 2004:3

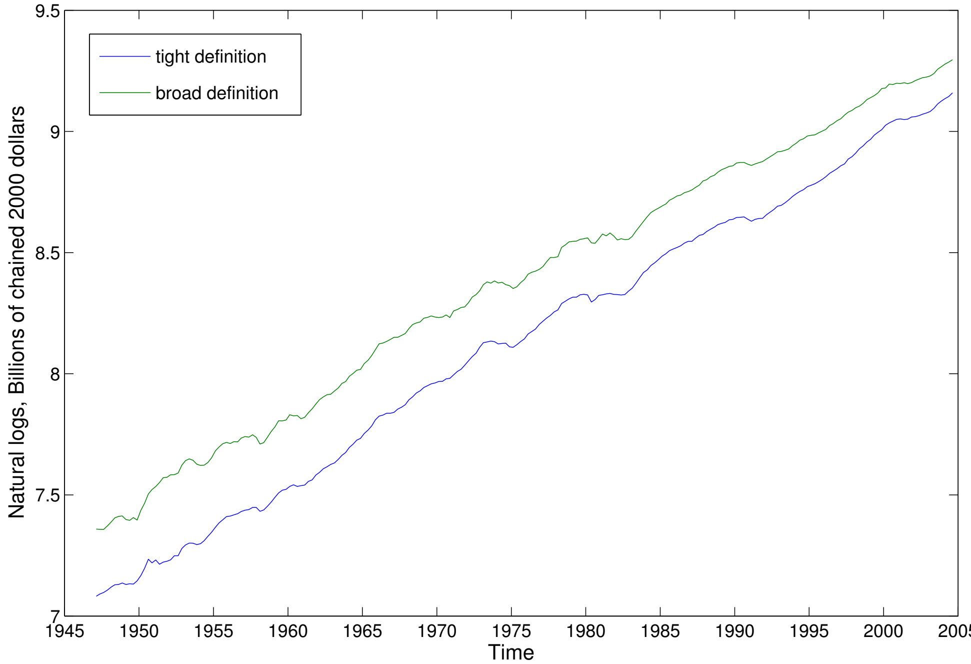

3.2 Our forecasting experiment

We estimate each forecasting model recursively and construct out-of-sample forecasts for each quarter from 1965q1 to 2001q3. Each quarter we:

1. Estimate all p,q combinations (where $p = 0, 1, 2, 3$ and $q = 0, 1, 2$), select best model using BIC.

2. Construct h-step ahead forecasts where h=1 to 8.

We compare the baseline ARMA model to our three theoretical models. These models replace step (1) above with:

1. model A: select ARMA(2,1) according to Campbell (1994), estimate ϕs and θs using maximum likelihood;

2. model B: select ARMA(2,1) according to Campbell (1994), use priors to estimate deep parameters to generate ϕs and θs using Bayesian methods;

3. model C: select ARMA(2,1) according to Campbell (1994), calibrate deep parameters that form ϕs and θs.

We collect the sets of pseudo real-time forecasts and compare them to the ex-post series detrended over the entire sample.[4]

For each of the forecast horizons (h), we define the 'loss', i ($\varepsilon_{i,t}$), from forecasting model i at time t to be the natural logarithm of the model's squared forecast error:

$$\varepsilon_{i,t} = \ln((\hat{y}^h_{i,t+h} - y^h_{t+h})^2) \tag{8}$$

We compute loss differentials (d_t) between model i and model j as $d_t = \varepsilon_{i,t} - \varepsilon_{j,t}$. This yields a series of loss differentials $\{d_t\}_{t=1}^{T}$, where $T = ((T_2 - h) - T_1)$ and T_1 and $T_2 - h$ are respectively the first and last dates over which the out-of-sample forecasts are made.

We test the null hypothesis that model i and model j have equal forecast accuracy using the test statistic due to Diebold and Mariano (1995):

$$TS = \frac{\bar{d}_t}{\sqrt{\hat{V}(\bar{d}_t)}} \tag{9}$$

where:

$$\bar{d} = \frac{100}{T} \sum_{t=1}^{T} d_t \tag{10}$$

and $\hat{V}(\bar{d})$ is an estimate of the asymptotic variance of $\bar{d}$, and $\bar{d}$ is the mean percentage difference in loss between models i and j.

We estimate $\hat{V}(\bar{d})$ using the Newey and West (1987) Heteroskedasticity and Autocorrelation Consistent (HAC) variance estimator, with a truncation lag of $(h-1)$. We improve small sample properties by making the following adjustment suggested by Harvey, Leybourne, and Newbold (1997):

$$TS^* = \sqrt{\frac{T^2 + T - 2Th + h(h-1)}{T^2}} TS \tag{11}$$

and compare this corrected test statistic to a Student-t distribution with $(T-1)$ degrees of freedom.

[4] The forecasts for the variables that are detrended using (log) differences are cumulated over the forecast period: a one-quarter-ahead forecast is a quarterly growth forecast, a two-quarter-ahead forecast is a forecast for growth over two quarters, and so on. These forecasts are then compared to the cumulative growth that occurred in the ex-post data.

3.3 Calibrating the Model

Campbell (1994) takes an agnostic view regarding the model parameterization, tabulating results across a range of permutations for ϕ=0.00,0.50,0.95,1.00 and σ=0,0.2,1,5,∞. Rather than make a statement regarding the plausibility of particular values, Campbell's (1994) strategy is to explore the implications of different parameter values. These are often special cases of other models in the literature (for example, under the fixed labour model if $\sigma = 0$ Campbell (1994) states his model is a general equilibrium version of Hall's (1978) permanent income model) or special cases in terms of agent's behaviour (for example, if $\sigma = \infty$, the agent is risk-neutral).

Our goal is not to explore the theoretical implications of alternative calibrations of Campbell's (1994) model. Instead, we seek to make quantitative inference statements about the model's forecasting ability. Based on other macroeconomic studies (for example, Cooley and Prescott (1995) and King and Rebelo (2000)) we regard certain parameter values as more likely than other values and calibrate Campbell's (1994) model accordingly.[5] These prior distributions are used as inputs for the Bayesian estimated model, and the means of the distributions are used as the parameter values for the calibrated model. Table 1 gives both our priors and posterior estimates for the model estimated over the longest sample (1947q1 to 2002q3) and detrending the narrow output definition with a linear trend.

Table 1: Prior and posterior distributions for Campbell (1994) model

	Prior				Posterior			
	Mean	Density	Range	SD	Mean	Median	CI(low)	CI(high)
ϕ	0.9	Beta	[0,1)	0.1	0.956	0.979	0.942	0.994
σ	1	Gamma	[0,∞)	0.2	0.345	0.619	0.150	0.956
σ_a	0.01	inv Gam	[0,∞)	4	0.012	0.012	0.010	0.013
α	0.667	Point		-	-	-	-	-
δ	0.025	Point	-	-	-	-	-	-
g	0.005	Point	-	-	-	-	-	-
r	0.015	Point	-	-	-	-	-	-

NB: The Inverse Gamma priors are of the form $p(\sigma|\nu, s) \propto \sigma^{-\nu-1}e^{-s/2\sigma^2}$ and we report the parameters s and ν.

Figures 2 and 3 show prior and posterior distributions for ϕ, the auto-

[5]That particular parameter values are more plausible than others is almost implicit in Campbell's (1994) exploration of the model's implication for highly autocorrelated ($\phi = 0$) technology shocks.

correlation in technology shocks, and σ, the intertemporal substitution of labour. The data shift the prior for ϕ towards a relatively tight prior, centred on highly correlated technology shocks. The data shift the mass of the prior for σ towards much lower values.

Figure 2: Indicative Campbell (1994) model: ϕ

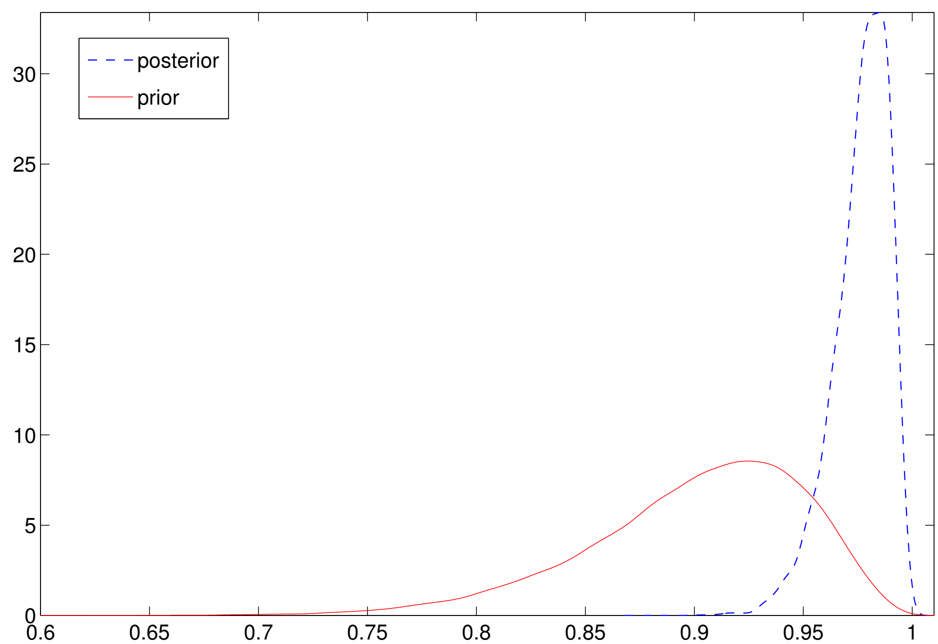

We estimate the posterior distribution of the parameters by combining our prior distributions with the likelihood function, estimated using the Kalman filter. We generate 50,000 draws from joint density of the parameters using the Metropolis-Hastings posterior simulator. The first 25,000 of draws from the posterior density are discarded, and we use the posterior means of the remaining draws as our estimates of the model parameters. Our pseudo real-time forecasting exercise requires us to re-estimate the model each quarter — for *every* combination of data definition and detrending method.[6] This makes the process very computationally intensive.

[6] We do not consider data revisions.

Figure 3: Indicative Campbell (1994) model: σ

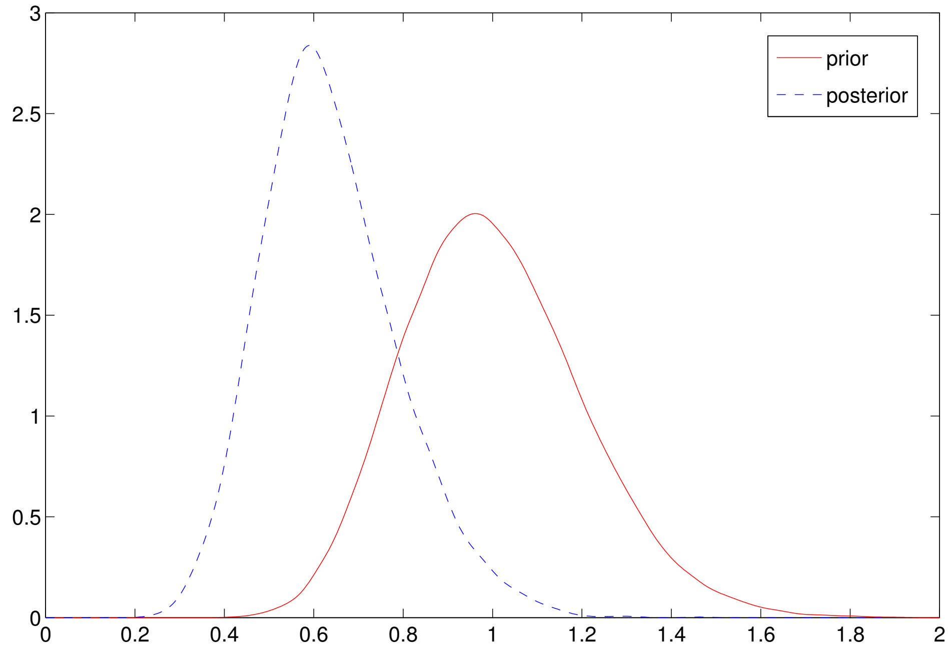

3.4 Forecast comparisons

Table 3 details the US output forecast comparisons between the baseline statistical ARMA model and the three theoretical models (A, B and C).[7] The table presents results for each data definition and each de-trending procedure. Negative numbers represent a lower mean differential loss from the theoretical model and hence a forecasting improvement for the theoretical model.

Focusing initially on the results for the narrow data definition (in the top half of the table), model A, that uses the Campbell (1994) model to inform the choice of ARMA order, performs particularly well. Detrending the data using an ex-post linear trend, model A outperforms the statistical model for all but three forecast horizons (see the first column labelled A in table 2). The six quarter ahead forecast is statistically significant (at the 10% level) but summing the loss differentials across all horizons does not yield a statistically significant forecast improvement. Improvements in forecast performance are

[7]Additional tests show that the ARMA model is a reasonable statistical test — ARMA forecasts yield lower mean square errors than a four variable VAR (using output, consumption, capital and labour) with four lags. Stock and Watson (2001) state that only “state-of-the-art VAR” models beat simple univariate ARMA models.

also realized when first differences are used to detrend the data. Column four of the table shows forecast improvement in seven of the eight horizons.

The performance of the Bayesian model (listed in columns labelled B in the table) is comparable with the statistical ARMA model.[8] Under the linear trend, model B performs relatively poorly at shorter forecast horizons but yields forecasting improvements at longer horizons. In contrast, when the data are first differenced, the initial five quarters yield improvements in forecasting with model B over the statistical ARMA model. At the third and fourth quarter in particular, mean improvements of 32% and 37% appear non-trivial. However, the performance of the model deteriorates at longer horizons. Still, compared with the statistical ARMA model — explicitly designed to forecast well — the Bayesian model appears at least competitive under both detrending methods.

Finally, the model calibrated according to the means of our priors performs particularly poorly and results in marked reduction in forecasting ability. A two-sided test of the null that the forecasting performance is not equal across models would entail doubling the p-values associated with each cell in the table. The cells in row "*sum*" of the top half of the table show we can reject the null that the forecasting performance of model C is the same as the statistical model — it is much worse.

The results in the second half of the table show a marked deterioration in performance for the theoretical models relative to results for the narrow definition of output in the first half of the table. Model A returns a statistically significant inferior performance at the first quarter. However, the remaining forecast horizons show better performance and in fact, selecting the ARMA order according to the theoretical model yields significantly better performance at the five quarter horizon. Results for this model using first differences are mixed.

Using linearly detrended data model B and model C show some forecasting improvements at longer horizons but these results are also mixed. Both models perform badly under the broad data definition using first differenced data: no quarter yields forecasts superior to the statistical ARMA model.

However, this deterioration in forecasting performance using a broad output measure is to be expected. Campbell's (1994) model is a model of consumption and investment, not the aggregate of output our broad measure represents. Quantitatively, thinking about the relevant dataset for theoretical model seems to matter. Richer models that treat explicitly each element of output may improve output forecasts from ARMA models.

[8]Results available on request show the Bayesian model beats a random walk at standard significance levels.

Table 2: Loss differentials for the United States

			Narrow			
		LT			FD	
h	model A	model B	model C	model A	model B	model C
1	-5.109	11.576	30.771	-11.807	-6.085	20.319
	(0.579)	(0.464)	(0.131)	(0.423)	(0.775)	(0.368)
2	-6.002	50.506	58.412	-13.111	-6.456	27.531
	(0.523)	(0.007)	(0.000)	(0.230)	(0.717)	(0.183)
3	6.114	9.793	39.044	-4.628	-32.390	-8.262
	(0.435)	(0.642)	(0.025)	(0.642)	(0.262)	(0.705)
4	-13.385	11.018	36.706	-11.829	-37.384	-18.355
	(0.198)	(0.582)	(0.030)	(0.308)	(0.140)	(0.448)
5	-0.048	-4.945	19.492	-9.209	-4.373	17.272
	(0.992)	(0.807)	(0.272)	(0.321)	(0.857)	(0.347)
6	**-15.723**	-9.114	10.691	4.498	27.616	47.557
	(0.078)	(0.666)	(0.489)	(0.524)	(0.249)	(0.035)
7	-5.241	-18.740	5.749	-1.178	36.589	60.377
	(0.447)	(0.474)	(0.618)	(0.889)	(0.143)	(0.029)
8	-5.930	-18.133	12.881	-9.718	44.957	29.606
	(0.64)	(0.568)	(0.225)	(0.109)	(0.083)	(0.229)
sum	-2.094	1.605	17.918	-0.487	15.648	30.565
	(0.414)	(0.920)	(0.043)	(0.875)	(0.400)	(0.045)
			Broad			
		LT			FD	
h	model A	model B	model C	model A	model B	model C
1	12.342	0.703	28.241	-2.796	23.817	34.413
	(0.014)	(0.968)	(0.070)	(0.824)	(0.130)	(0.040)
2	-0.355	17.120	14.701	14.805	26.243	62.821
	(0.965)	(0.194)	(0.361)	(0.327)	(0.203)	(0.014)
3	-5.566	8.234	3.790	-2.197	19.168	35.575
	(0.299)	(0.622)	(0.775)	(0.875)	(0.365)	(0.022)
4	-5.042	-15.482	-2.069	8.088	17.877	29.770
	(0.426)	(0.495)	(0.879)	(0.623)	(0.329)	(0.165)
5	**-14.877†**	-8.742	8.629	17.916	39.580	68.620
	(0.047)	(0.718)	(0.378)	(0.263)	(0.027)	(0.005)
6	-10.403	-14.678	-6.811	27.457	40.992	76.150
	(0.233)	(0.546)	(0.576)	(0.175)	(0.158)	(0.005)
7	5.589	-13.895	-1.399	20.338	24.837	93.765
	(0.445)	(0.638)	(0.889)	(0.199)	(0.329)	(0.001)
8	-0.241	11.097	8.083	40.802	49.254	106.754
	(0.985)	(0.691)	(0.387)	(0.048)	(0.085)	(0.002)
sum	-0.603	-3.808	6.085	10.898	22.559	64.632
	(0.792)	(0.806)	(0.176)	(0.278)	(0.109)	(0.001)

NB: * 10% sig. level; †5% sig. level; ‡1% sig. level, one-sided test.
sum is the mean loss differential cumulating the squared errors over all horizons.

4 Conclusion

Our results show that information from a parsimonious RBC model can improve forecasts from ARMA models by selecting the order of the ARMA process. In addition, bringing theoretically derived prior information to the data yields forecasting performance comparable to that of the statistical ARMA model that is explicitly designed to forecast well.

These results are sensitive to the definition of output the model attempts to forecast. The model is capable at forecasting a narrow definition of output only, a richer model that seeks to explain additional components of output may yield similar improvements. Future work could explore the implications of the RBC model for forecasting consumption and the other variables contained in the Campbell (1994) model.

References

Box, G E P and G M Jenkins (1976), *Time Series Analysis: Forecasting and Control*, Holden-Day, San Francisco.

Campbell, J Y and S Ludvigson (2001), "Elasticities of substitution in real business cycle models with home production," *Journal of Money, Credit and Banking*, 33(4), 847–875.

Canova, F (1998), "Detrending and business cycle facts," *Journal of Monetary Economics*, 31(3), 475–512.

Chari, V V, P J Kehoe, and E R McGrattan (2005), "A critique of structural VARs using real business cycle theory," (631).

Clements, M P and D F Hendry (1998), *Forecasting Economic Time Series*, Cambridge University Press, Cambridge, U.K.

Cooley, T F and E C Prescott (1995), "Economic growth and business cycles," chap 1.

DeJong, D N, B F Ingram, and C H Whiteman (2000), "A Bayesian approach to dynamic macroeconomics," *Journal of Econometrics*, 98(2), 203–23.

Del Negro, M and F Schorfheide (2004), "Priors from general equilibrium models for VARs," *International Economic Review*, 45(2), 643–73.

Diebold, F X and R S Mariano (1995), "Comparing predictive accuracy," *Journal of Business and Economic Statistics*, 13(3).

Enders, W (1995), *Applied econometric time series*, Wiley series in probability and mathematical statistics., Wiley, New York, walter Enders. ill. ; 24 cm.

Fernandez-Villaverde, J, J Rubio-Ramirez, and T Sargent (2005), "A,B,C's (and D)'s for understanding VARs," .

Greenwood, J, Z Hercowitz, and G Huffman (1988), "Investment, capacity utilization, and the real business cycle," *American Economic Review*, 78(3), 402–417.

Hamilton, J D (1994), *Time series analysis*, Princeton University Press, Princeton, N.J., James D. Hamilton. ill. ; 26 cm.

Harvey, D S, S Leybourne, and P Newbold (1997), "Testing the equality of prediction Mean Squared Errors," *International Journal of Forecasting*, 13, 117–31.

Ingram, B F and C H Whiteman (1994), "Supplanting the 'Minnesota' prior: Forecasting macroeconomic time series using real business cycle model priors," *Journal of Monetary Economics*, 34(3), 497–510.

Kapetanios, G, A Pagan, and A Scott (2005), "Making a match: Combining theory and evidence in policy-oriented macroeconomic modelling," *Centre For Applied Macroeconomic Analysis Working Paper 2005/1*, The Australian National University.

King, R G and S T Rebelo (2000), "Resuscitating real business cycles," *NBER Working Paper*, (7534).

Lettau, M (2003), "Inspecting the mechanism: Closed-form solutions for asset prices in real business cycle models," *Economic Journal*, 15, 550–75.

Ludvigson, S (1996), "The macroeconomic effects of government debt in a stochastic growth model," *Journal of Monetary Economics*, 38, 25–45.

Newey, W K and K West (1987), "A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55, 703–708.

Stock, J H and M W Watson (2001), "Vector autoregressions," *Journal of Economic Perspectives*, 15(4), 101–15.

5 Appendix

The gross rate of return on a one-period investment in capital R_{t+1} is defined to be the marginal product of capital plus the capital stock without depreciation:

$$R_{t+1} \equiv (1-\alpha)\left(\frac{A_{t+1}N_{t+1}}{K_{t+1}}\right)^{\alpha} + (1-\delta) \tag{12}$$

The first order conditions that result from maximising the objective function (3) subject to the constraints (4) and (5) are:

$$C_t^{-\gamma} = \beta E_t\{C_{t+1}^{-\gamma} R_{t+1}\} \tag{13}$$

and:

$$\theta(1-N_t)^{-\gamma} = \frac{W_t}{C_t} = \alpha\frac{A_t^{\alpha}}{C_t}\left(\frac{K_t}{N_t}\right)^{1-\alpha} \tag{14}$$

where the marginal utility of leisure is set equal to the wage W_t times the marginal utility of consumption. The wage also equals the marginal product of labour from the production function.
In steady-state, technology, capital, output, and consumption all grow at a common constant rate G, where $G \equiv A_{t+1}/A_t$. In the steady state, the gross rate of return on capital R_{t+1} is constant R, so that the first-order condition (13) becomes:

$$G^{\gamma} = \beta R \tag{15}$$

or in logs (denoted by lower-case letters),

$$g = \frac{\log(\beta) + r}{\gamma} = \sigma \log(\beta) + \sigma r. \tag{16}$$

The definition of the return to capital (12) and the first-order condition (13) imply that the technology-capital ratio is a constant:

$$\frac{A}{K} = \left[\frac{G^{\gamma}/\beta - (1-\delta)}{1-\alpha}\right]^{1/\alpha} \approx \left[\frac{r+\delta}{1-\alpha}\right]^{1/\alpha} \tag{17}$$

using $G^{\gamma}/\beta = R \approx 1+r$. The production function and the technology-capital ratio (17) also imply a constant steady-state output-capital ratio:

$$\frac{Y}{K} = \left(\frac{A}{K}\right)^{\alpha} \approx \frac{r+\delta}{1-\alpha} \tag{18}$$

The steady state consumption-output ratio is then a constant given by:

$$\frac{C}{Y} = \frac{C/K}{Y/K} \approx 1 - \frac{(1-\alpha)(g+\delta)}{r+\delta} \tag{19}$$

When the model is outside its steady state it is a system of non-linear equations, in logs of technology, capital, output, and consumption, caused by time variation in the consumption-output ratio and by incomplete capital depreciation. Campbell (1994) obtains an approximate analytical solution for the model by transforming the model into a system of approximate log-linear difference equations and taking first-order Taylor series approximations around the steady state:

$$y_t = \alpha(a_t + n_t) + (1-\alpha)k_t \tag{20}$$

$$k_{t+1} \approx \lambda_1 k_t + \lambda_2(a_t + n_t) + (1 - \lambda_1 - \lambda_2)c_t \tag{21}$$

$$r_{t+1} \quad = \quad \lambda_3(a_{t+1} + n_{t+1} - k_{t+1}) \tag{22}$$

$$E_t \triangle c_{t+1} = \lambda_3 E_t[a_{t+1} + n_{t+1} - k_{t+1}] \tag{23}$$

$$n_t = \nu[(1-\alpha)k_t + \alpha a_t - c_t] \tag{24}$$

To close the model, technology is assumed to follow an AR(1) process:

$$a_t = \phi a_{t-1} + \epsilon_t \tag{25}$$

Campbell solves the model analytically using the method of undermined coefficients, and shows that the dynamic behavior of the economy can be characterised by the following equations.

$$c_t \equiv \eta_{ck} k_t + \eta_{ca} a_t \tag{26}$$

$$n_t \equiv \eta_{nk} k_t + \eta_{na} a_t \tag{27}$$

$$y_t \equiv \eta_{yk} k_t + \eta_{ya} a_t \tag{28}$$

$$k_{t+1} \equiv \eta_{kk} k_t + \eta_{ka} a_t \tag{29}$$

where:

$$\eta_{ck} \equiv \frac{-Q_1 - \sqrt{Q_1^2 - 4Q_0 Q_2}}{2Q_2}$$

$$Q_0 \equiv -\lambda_3((1-\alpha)\nu - 1)(\lambda_1 + \lambda_2(1-\alpha)\nu)$$

$$Q_1 \equiv (1+\lambda_3\nu)(\lambda_1 + \lambda_2(1-\alpha)\nu) - \lambda_3((1-\alpha)\nu - 1)(1-\lambda_1-\lambda_2(1+\nu)) - 1$$

$$Q_2 \equiv (1+\lambda_3\nu)(1-\lambda_1-\lambda_2(1+\nu))$$

$$\eta_{ca} \equiv \frac{(1+\alpha\nu)(\lambda_3\phi - \lambda_2(\eta_{ck}(1+\lambda_3\nu) - \lambda_3((1-\alpha)\nu-1)))}{(\eta_{ck}(1+\lambda_3\nu) - \lambda_3((1-\alpha)\nu-1))(1-\lambda_1-\lambda_2(1+\nu)) - (1-\phi(1+\lambda_3\nu))}$$

$$\eta_{nk} \equiv \nu(1-\alpha-\eta_{ck})$$

$$\eta_{na} \equiv \nu(\alpha-\eta_{ca})$$

$$\eta_{kk} \equiv \lambda_1 + \lambda_2(1-\alpha)\nu + \eta_{ck}[1-\lambda_1-\lambda_2(1+\nu)]$$

$$\eta_{ka} \equiv \lambda_2(1+\alpha\nu) + \eta_{ca}[1-\lambda_1-\lambda_2(1+\nu)]$$

$$\eta_{yk} \equiv (1-\alpha) + \alpha\nu(1-\alpha-\eta_{ck})$$

$$\eta_{ya} \equiv \alpha + \alpha\nu(\alpha-\eta_{ca})$$

$$\lambda_1 \equiv \frac{1+r}{1+g}, \quad \lambda_2 \equiv \frac{\alpha(r+\delta)}{(1-\alpha)(1+g)}, \quad \lambda_3 = \frac{\alpha(r+\delta)}{r+1}$$

$$\nu \equiv \frac{(1-N)\sigma}{N+(1-\alpha)(1-N)\sigma}$$