Technology Shocks or Coloured Noise? Why real-business-cycle models cannot explain actual business cycles

KEVIN D. HOOVER & KEVIN D. SALYER
Department of Economics, University of California, Davis, CA 95616-8578, USA

Typically real-business-cycle models are assessed by their ability to mimic the covariances and variances of actual business cycle data. Recently, however, advocates of RBC models have used them to fit the historical path of real GDP using the Solow residual as a driving process. We demonstrate that the success of RBC models at matching historical GDP data does not confirm the validity of RBC models. Through simulations we demonstrate that the Solow residual does not carry useful information about technology shocks and that the RBC model does not add incremental information about GDP. RBC models fit historical GDP data entirely because the Solow Residual is itself just a noisy measure of GDP.

1. Introduction

Macroeconometric models have traditionally been evaluated according to the accuracy with which they fit historical time-series data. In contrast, the proponents of real-business cycle models, following Kydland & Prescott (1982), have typically advocated a different standard of evaluation: models are judged on their ability to imitate, not the actual historical realizations of the stochastic processes that characterize the economy, but the typical operating characteristics of those processes as reflected in the relative variances and covariances or intercorrelations of business-cycle data. A real-business-cycle model is judged to be good, for example, not because it can fit the actual paths of consumption or investment, but because it reproduces, to nearly the correct magnitude, the greater smoothness of consumption relative to investment observed in the actual economy.

Despite this typical focus on stochastic operating characteristics, the lure of historical time-series has proved strong. Plosser (1989) appears to be the first real-business-cycle modeler to compare the path of output from a calibrated real-business-cycle model with the actual data. More recently, Hansen &

The authors would like to thank Edmund Snelling for able research assistance and Steven Sheffrin, Olivier Blanchard and the participants in the U.C. Davis Macroeconomics Workshop and seminars at U.C.L.A., University of California, Santa Barbara, the Federal Reserve Bank of San Francisco and the Institute for Advanced Studies, Vienna, for comments on earlier drafts.

0953-8259/98/030299-29

Prescott (1993) use a real-business-cycle model to explain the 1990–91 recession, while Thomas Cooley (1995) uses one to examine the causes of the 'Volcker recession' of 1979–82.[1]. In all these cases, the results appear plausible on the surface. Figure 1 presents a typical picture: it plots output from a real-business-cycle model against actual output (1964–93); the fit appears close.[2] The question is, why? Is it because the real-business-cycle models capture essential truths about the economy? Or is it merely a spurious correlation induced by the way in which the data are processed?

Advocates of real-business-cycle models believe of course that the good fit reflects the fact that the models capture some important truths about the economy. In these exercises, the real-business-cycle model is used to process information about technology shocks, as measured by the Solow residual, into series for output (and other macroeconomics aggregates). If the exercise is successful for good economic reasons, it is, first, because the Solow residual measures the relevant technology shocks; and, second, because the model itself adds information (based on its having captured some of the critical elements of the macroeconomic structure). We provide evidence that, with respect to modeling the path of output, both points are wrong: the Solow residual does not carry useful information about technology shocks; the real-business-cycle does not add incremental information.

Some might argue that, even if we are correct, our point is not relevant to the assessment of real-business-cycle models because those models are not 'about' realized historical paths, but about stochastic operating characteristics. It is worth reiterating that prominent advocates of real-business-cycle models themselves have recently used the models for the analysis of the realized path of output. It is therefore essential that we know whether these types of models can be legitimately put to such a use. At the same time we underscore the limited nature of our analysis: the ability to support conditional forecasts of a model is only one criterion on which to evaluate models; failure on this dimension does not imply that the model fails in all dimensions. We do not therefore pretend to present evidence against real-business-cycle models *per se*, but rather against the use of the Solow residual as a measure of the technology shocks that drive them and against the use of the models themselves to explain the historical path of output. The ability of the models to capture the historical behavior of other aggregates (e.g. consumption, investment or labor) conditional on output is beyond the scope of this paper and remains an open question.

Our conclusion will not come as a surprise in some circles. Hall (1986, 1990) and Hartley (1994a, b) have attacked the adequacy of the Solow residual as a measure of technology shocks. Watson (1993) and Cogley & Nason

[1] The recent focus by advocates of real-business-cycle models on comparing actual time series to conditional model forecasts may in part be due to an acknowledgment that such a methodology is in line with the Cowles Commission approach to econometric evaluation of models. This issue was highlighted in Hansen & Heckman's (1996) critique of calibration.

[2] The model on which this figure is based uses linear detrending and is described in the Appendix.

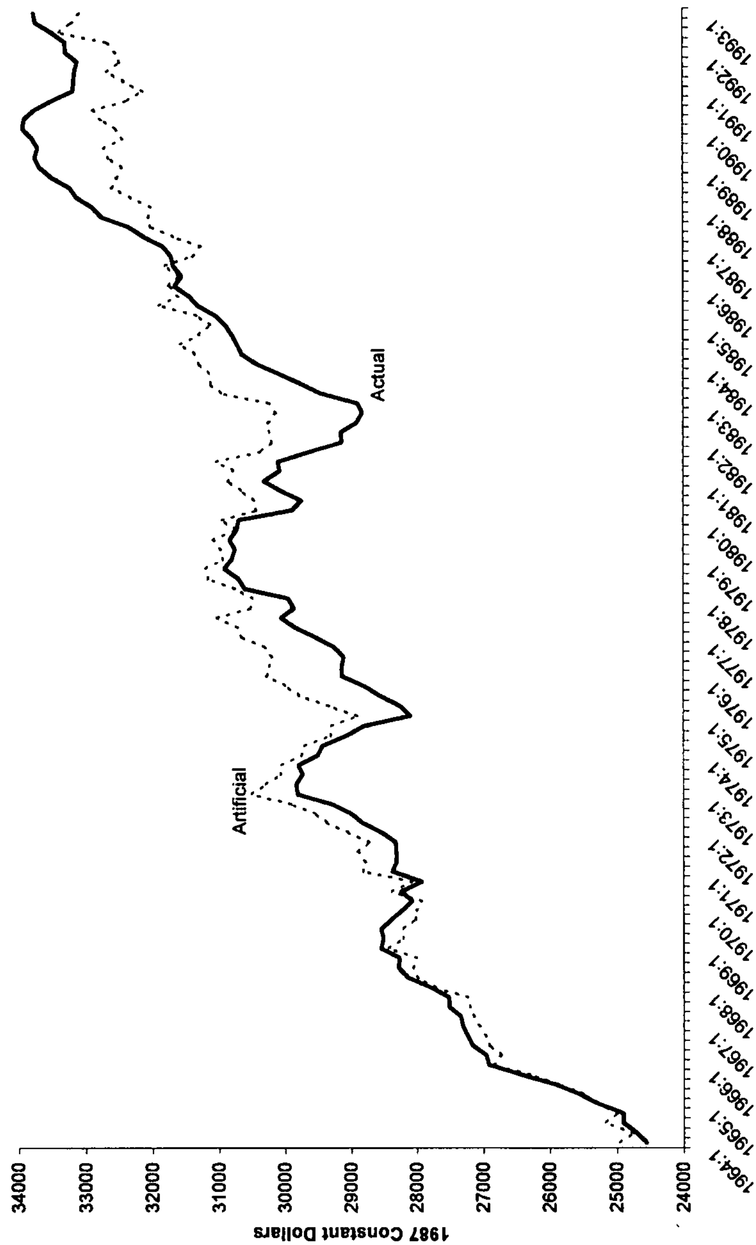

Fig. 1. Comparison of actual and artificial per-capita real GNP.

(1995b) provide evidence that the real-business-cycle model is not itself responsible for the cyclical characteristics of its simulated output. Canova (1991) and Cogley & Nason (1995a) provide evidence that some of the properties of output from real-business-cycle models are artifacts of the filtering procedures. The story we tell is distinguished from these earlier contributions in two respects. First, the previous work does not address the use of the real-business-cycle model in explaining particular historical episodes. Second, our evidence is presented in a readily comprehendable, heuristic way.

2. Taking Kydland & Prescott Seriously

Kydland & Prescott (1991, 1995) argue that real-business-cycle models systematically distort some aspects of economic reality in order to cast other aspects in higher relief. The models embody 'purposeful inconsistencies' in some dimensions so that they will be consistent in other dimensions (Kydland & Prescott, 1995, pp. 14–15). They argue that recognition of this feature calls into question the usefulness of traditional econometric estimation of the models—what they call the 'systems-of-equations approach'. One should not expect a calibrated real-business-cycle model to fit the data on ordinary statistical criteria as well as an econometrically estimated model: first, because the estimated model generally has numerous free parameters that are selected precisely to secure a good fit (Kydland & Prescott, 1991, p 170; Lucas, 1981, p. 288; Prescott, 1983, p. 11) and second because goodness-of-fit measures are typically symmetric and therefore penalize the purposeful inconsistencies that Kydland & Prescott believe to be essential to the power of real-business-cycle models to illuminate interesting questions (cf. Lucas, 1987, p. 45). Hoover (1995) argues that Kydland & Prescott's position has a legitimate methodological foundation, but points out the risk that their methods might make systematic empirical assessment of calibrated models impossible. In particular, since real-business-cycle modellers typical evaluate their models by a subjective comparison of selected sample moments of the output of their simulated models with the sample moments of the actual data, it is unclear on what neutral basis one might discriminate between two calibrated models founded on very different theoretical principles which nonetheless presented similar subjective matches to the actual data. It is not enough to say that '[t]he degree of confidence in the answer depends on the confidence that is placed in the economic theory being used' (Kydland & Prescott, 1991, p. 171) or that '... currently established theory dictates which [model] should be used' (Kydland & Prescott, 1991, p. 174), when the issue is the empirical basis on which economists should regard one theory as better established and, therefore, as more worthy of confidence than another.

Time-series econometrics permit us to characterize data in atheoretical ways. If real-business-cycle models carry any useful truths about the economy, they must provide us with some advantage relative to atheoretical time-series models. This advantage could cut in either of two directions. First, the predicted series for output from a real-business-cycle model might carry theoretically informed and interpretable information that permit us to rely on fewer atheoretical relationships in a relatively complete time-series characterization of the data.

We would therefore be able to reduce the number of free parameters needed to provide such a characterization.[3] Second, if we begin with an incomplete atheoretical characterization of the data, predicted output should provide additional information and help us to offer a more complete characterization.

The current research started in the first direction. We began with a vector autoregression in which quarterly output was regressed on 12 lags of itself and on the current and 12 lagged values of consumption, investment and hours of labor. We expected, with such a rich characterization of the data, that the predicted output of the real-business-cycle model would not turn out to be significant when it was added as an additional regressor. We intended to judge the success of the real-business-cycle model by posing the question, how much could we restrict the original VAR model plus the predicted output (i.e. how many free parameters could we eliminate from the atheoretical model) while maintaining the requirement that the more parsimonious form remain a statistically valid restriction of the most general specification.[4] When we added the predicted output from the real-business-cycle model to the VAR, contrary to our expectations, model output was not statistically insignificant, which would have indicated a redundant regressor, but highly significant. Furthermore, we could reject the hypothesis of no serial correlation at very high levels of significance, even though the VAR without predicted output as a regressor appeared to have white-noise errors.

It turns out, and this is the version that we build on for the rest of this paper, that the exact same phenomena occur when predicted output is added to a univariate time-series model: that is to say, predicted output from the real-business-cycle model has incremental predictive power for actual output even after the past history of actual output is taken into account. This is indeed the econometric analogue of Fig. 1. And it raises the same question: why? The obvious answer, and surely the one that real-business-cycle modelers would prefer, is that the model is highly explanatory for output. The remainder of this paper is devoted to showing that this explanation is incorrect.

To test whether the model adds information, we use the Solow residual itself as a regressor in the output equation (actual output, actual Solow residual). If it were to dominate model output, then the model would be shown to have had no value added. In fact, it does dominate model output in the sense of having a higher level of significance in the regression. This is, however, muddied by the fact that model output remains independently significant: dominance is incomplete.

We then construct artificial data for the Solow residual using actual data for output but artificial data with similar stochastic properties to the actual factor input series to construct a *faux* Solow residual. Since this *faux* Solow residual is just noisy output (of a particular coloration), we know in advance that the model cannot have any value added, because there are, by construction, no technology

---

[3] New classicals, and especially advocates of calibrated business-cycle models, hold the elimination of free parameters to be the chief virtue of their preferred methods (see Lucas, 1980, p. 288; Prescott, 1983, p. 11).

[4] A fuller description of the proposed evaluation method and its rationale are found in Hoover (1994).

shocks to process. We find that, even in this rigged case, the artificial data reproduces the relationship between the Solow residual and the model output that we found with actual data (i.e. the fact that dominance is incomplete). This removes any reason for regarding incomplete dominance as evidence for value added on the part of the real-business-cycle model.

We regard this last test as bearing on the question of whether the Solow residual measures technology shocks at the business-cycle frequency as well. There is no independent evidence that the Solow residual captures technology shocks. The best evidence is the supposed success of the real-business-cycle model driven by such shocks in matching the actual data. To the extent that this evidence is just the relative volatilities of the various time-series generated by repeated simulation of the real-business-cycle model compared to the actual relative volatilities (a standard type of evidence in the real-business-cycle literature) our paper is silent. But real-business-cycle proponents have claimed success in fitting the actual time series for output as well. The test described in the previous paragraph shows that a model driven by an artificial Solow residual which, by construction, does not capture technology shocks, behaves statistically just like a real-business-cycle driven by the actual Solow residual. Thus, the goodness of fit between the model output and the actual output does not appear to depend on the Solow residual actually capturing technology shocks, but on other statistical and model characteristics independent of true technology shocks.

To reinforce our conclusion, we repeat these tests with a Solow residual constructed with an artificial series for output as well as for the factors of production. The fact that identical statistical properties are reproduced in the completely artificial case (a case in which there is, by construction, no linkage between artificial output and artificial factors of production or technology) demonstrates that none of these statistical properties can be taken as evidence in favor of the actual Solow residual having captured actual technology shocks at business-cycle frequencies, for they appear even when, by construction, there are no technology shocks to capture.

Up to this point, our results are based on linearly filtered data. In a final exercise, we highlight the importance of filtering procedures by showing that results based on the Hodrick–Prescott filter—the most common detrending procedure in the real-business-cycle literature—even more dramatically illustrate the artifactual nature of predicted output from real-business-cycle models.

3. Model, Data, Calibration

We begin with a model similar to one used by Hansen (1985). Like all competitive-equilibrium real-business-cycle models it is a dynamic optimization model for a representative agent.[5] The model is fundamentally a neoclassical growth model, but in typical fashion it has been cast in such a manner as to explain deviations from steady-state growth paths. The model is calibrated by choosing values for its key parameters on the basis of national accounting

[5] McCallum (1989) provides a particularly clear introduction to real-business-cycle models.

Table 1. Descriptive statistics for US economy: 1964:1–1993:1 and artificial economy (actual data have been detrended using a linear time trend, model data is unfiltered

	% SD of output	σ_c/σ_y	σ_i/σ_y	σ_h/σ_y	*Corr(y,c)*	*Corr(y,i)*	*Corr(y,h)*
US data	3.49	0.90	2.23	0.73	0.93	0.68	0.61
Model	3.72	0.67	3.02	0.45	0.81	0.88	0.76
King *et al.* Model*	4.26	0.64	2.31	0.48	0.76	0.85	0.73

*Taken from King *et al.* (1988), Table 5, p. 224. For their analysis, the autoregressive parameter for the technology shock, ρ, was assumed to be 0.90 rather than the value of 0.96 which we use.

considerations or independent (microeconomic) studies. The data used in the calibration exercise and in our evaluation of the model are constructed following Cooley & Prescott (1994). (The model is described in detail, along with its calibration and the construction of the data in the Appendix.)

The source of aggregate fluctuations is technology shocks. Technology shocks are measured by the Solow residual (Solow, 1957):

$$SR = Y - \alpha K - (1 - \alpha)L \tag{1}$$

where Y, L and K are the logarithms of GNP, labor and capital, and α is capital's share of output (precise definitions are in the Appendix).[6]

Typically, real-business-cycle models have been assessed by estimating a univariate time-series model of the Solow residual in order to obtain a good characterization of its dynamic properties. The models are then simulated numerous times drawing the technology shocks from a distribution with those same properties. The variances and covariances of the endogenous variables of the model (in the case of the current model, output, consumption, investment and hours) are then compared informally with the second moments of the actual detrended series. Table 1 compares key variances and covariances of the output of the model with the actual data. The behavior of the relative (to GNP) volatilities and contemporaneous correlations with GNP of the endogenous variables in the model we study are similar to those of the model of King *et al.* (1988)[7]. In some dimensions, our model is somewhat closer to the actual data; in others, their model has the edge; in all, the differences are small. Both models duplicate some of the key features of the data: consumption is less volatile than output, which is less volatile than investment. Also, the correlation of both consumption and investment with output in the models is roughly in line with that in the data. However, labor is too highly correlated with output in both models while, at the same time, both models underpredict the volatility of labor

[6] It is also common to see the Solow residual calculated as a percentage change: $\Delta SR = \Delta Y - \alpha\Delta K - (1 - \alpha)\Delta L$.

[7] We choose this model as a fairly standard real-business-cycle model that used a linear detrending procedure like the one we use.

Table 2. Regression results, linearly detrended output

Line	Coefficient estimates[1]: Additional variables	Coefficient (t-statistics)	Summary statistics: $\bar{R}^2$	SER^2	F statistic[3]	Regression diagnostics: Normality[4]	Serial correlation up to order:[5] 1	2	4	Arch[6]	Hetero-skedasticity[7]
1	None		0.95	0.0076	293.39	10.52	0.38	1.21	1.02	1.28	1.46
							(0.55)	(0.30)	(0.40)	(0.28)	(0.13)
							0.40	2.61	4.40	5.10	22.16
							(0.53)	(0.27)	(0.35)	(0.28)	(0.14)
2	*YM*	0.14 (4.68)	0.96	0.0069	313.77	11.61	101.92	50.77	25.04	1.44	1.04
							(0.00)	(0.00)	(0.00)	(0.22)	(0.42)
							57.35	57.52	57.69	5.74	18.83
							(0.00)	(0.00)	(0.00)	(0.22)	(0.40)
3	*SRL*	0.31 (6.26)	0.96	0.0065	357.41	10.46	87.16	43.63	21.81	1.19	0.92
							(0.00)	(0.00)	(0.00)	(0.32)	(0.56)
							52.80	53.10	53.65	4.77	16.86
							(0.00)	(0.00)	(0.00)	(0.31)	(0.53)
4	*YM*	− 0.44 (4.37)	0.97	0.0060	378.12	11.90	54.42	28.30	14.76	0.66	0.94
							(0.00)	(0.00)	(0.00)	(0.62)	(0.54)
	SRL	1.04 (5.99)					39.94	41.23	42.79	2.71	19.21
							(0.00)	(0.00)	(0.00)	(0.61)	(0.51)
5	*FYM*	0.08 (3.22)	0.96	0.0057	284.60	7.35	248.74	123.21	60.50	1.33	1.04
							(0.00)	(0.00)	(0.00)	(0.26)	(0.43)
							82.04	82.04	82.07	5.31	18.73
							(0.00)	(0.00)	(0.00)	(0.26)	(0.41)

6	*FSRL*	0.22 (4.97)	0.96	0.0050	320.73	4.59	233.67	121.37	59.79	1.12	0.84
							(0.00)	(0.00)	(0.00)	(0.35)	(0.64)
							80.49	81.67	81.78	4.48	15.76
							(0.00)	(0.00)	(0.00)	(0.34)	(0.61)
7	*FYM*	− 0.53 (6.50)	0.97	0.0058	404.09	1.35	118.32	60.64	30.61	0.65	1.04
							(0.00)	(0.00)	(0.00)	(0.63)	(0.42)
	FSRL	1.15 (7.76)					61.99	62.99	63.82	2.67	20.87
							(0.00)	(0.00)	(0.00)	(0.62)	(0.40)

Each line reports a regression of *YL* on a constant and eight own lags and additional variables using quarterly data, 1964:1–1993:1.

1 Coefficient estimates are not reported for lags of *YL* or the constant.

2 Standard error of regression.

3 Tests the null hypothesis that coefficients of all regressors are zero.

4 Jarque–Bera test for normality of residuals.

5 Breusch–Godfrey tests for serial correlation. The top pair of numbers is the F-form. The bottom pair is (number of observations) $\times R^2$ form which is distributed asymptotically as Chi-square. *P*-values are given in parentheses.

6 Tests the null of no autoregressive conditional heteroskedasticity. The top pair of numbers is the F-form. The bottom pair is the (number of observations) $\times R^2$ form which is distributed asymptotically as Chi-square. *P*-values are given in parentheses.

7 Tests the null of no heteroskedasticity. The top pair of numbers is the F-form. The bottom pair is the (number of observations) $\times R^2$ form which is distributed asymptotically as Chi-square. *P*-values are given in parentheses.

relative to GNP. The inability of the basic model to duplicate the features of the labor market motivates several modifications to the model, e.g. indivisible labor (Hansen, 1985), including government purchases (Christiano & Eichenbaum, 1992), and including a household production sector (Benhabib, *et al.* 1991).

To use real-business-cycle models to explain actual historical episodes: (1) the Solow residual (*SR*) is calculated for the whole sample and then linearly detrended (we call the detrended Solow residuals *SRL*)[8]; (2) *SRL* is used as the driving process for the calibrated model to generate model output (i.e. the calculated values of the endogenous variables from the model conditional on *SRL*); (3) the detrended actual data and the model data can then be compared either directly (as we shall do later in the paper) or by adding back the trends extracted from the actual data (as reflected in Fig. 1).

4. The Puzzle

In Section 2 above, we noted that we had expected that predicted output would not prove to be a significant regressor when added to a richly parameterized VAR. It not only turned out to be highly significant, but it also appeared to induce serial correlation. We find this result puzzling. We also discovered that the same phenomena are present when we add predicted output to a univariate time-series model for actual output. Its statistical significance is perhaps less puzzling than it was with the richer VAR, but we are still entitled to ask whence comes its incremental predictive power—from the fact that the real-business-cycle model captures some essential truth about the economy or from some artifact of the data? The strange induction of serial correlation may suggest the latter. To keep things simple, we shall investigate the competing explanations using the univariate time-series models as the base regressions.

Consider actual output linearly detrended (*YL*) A good description of the stochastic properties of this series can be found in the following univariate time-series model:

$$YL_t = \beta_0 + \sum_{j=1}^{8} \beta_j YL_{t-j} + \varepsilon_t \tag{2}$$

An estimate of Equation (2) is reported in Table 2, line 1. The estimates of the coefficients, β_j, are of little interest and are not reported, but the fit of the equation as a whole and the properties of the error term tell us that Equation (2) describes *YL* reasonably well. One should note the relatively high explanatory power of the regression as measured by $\bar{R}^2$, standard error of regression (SER) and the significance of the regressors as a group (F-statistic). Most important, one should note that there is reasonable evidence that the residuals are white-noise. Observe especially that there is no evidence of serial correlation up to fourth order at conventional significance levels.

[8] A suffix L on a variable name indicates that it is linearly detrended; a suffix HP indicates that it is detrended using the H-P filter.

Now consider what happens when we add the prediction of output from our real-business-cycle model to the regression. We now estimate:

$$YL_t = \beta_0 + \sum_{j=1}^{8} \beta_j YL_{t-j} + \gamma YM_t + \varepsilon_t \tag{3}$$

where *YM* is the prediction of output from the model.

One night expect that since Equation (2) is so richly parameterized that *YM* would add little explanatory power to the regression—indeed that was our initial conjecture. Such a result would not mean that the model is a bad one. Rather, one might expect that the model could be interpreted as imposing restrictions on the univariate time-series model, so that in the face of a profligately parameterized model such as Equation (2) the model data would appear redundant.

As we observed at the outset, in practice we find something quite different. Table 2, line 2 shows that our initial conjecture is wrong: the standard error of regression falls by 10% from 0.0076 to 0.0069 and the t-statistic on γ, the coefficient on *YM*, is highly significant.[9]

At first sight, this might appear to be overwhelming evidence that the model has explanatory power in addition to the past history of *YL* itself. We, however, prefer an alternative explanation: *YM* is just a noisy measure of *YL* and Equation (3) fits so well because it regresses the dependent variable on itself.

One initial piece of evidence for this explanation is found in the tests of serial correlation. For serial correlation up to fourth order the Breusch–Godfrey test statistics are $F = 25.04$ (p-value $= 0.00$) and $\chi^2 = 57.69$ (p-value $= 0.00$), which reject the null of insignificant serial correlation at any conventional level. (The table presents similar results for serial correlation up to first and second orders.) The introduction of *YM* as a regressor not only raises the fit of the regression, but also induces serial correlation where there was no evidence of it before.

This is an odd situation. One normally thinks that additional regressors, provided that there are no problems with degrees of freedom, do not degrade the regression. One possibility is the following: *YM* is simply a noisy measure of *YL*. When it is introduced into the regression, it is so well correlated with *YL* that it massively increases the precision of the fit. Having better tuned the signal from *YL*, it then permits serial correlation that was present in the data, but which itself presented a weak signal, to rise above the noise.[10]

---

[9] It is worth reminding the reader of the point that we made at the outset of this section and in the introduction that this same phenomena also occurs in a richly parameterized VAR that includes current information on consumption, investment, and labor hours. It is, therefore, not the case that the puzzle somehow arises only because model output carries current information that would not otherwise be available. In any case, even if that were so, so that one found the puzzle less puzzling than we do ourselves, the question remains open as to what is the nature of this current information: is it information about technology shocks from the Solow residual with economic information added by the model? Or is it nothing more than information about output itself? We believe that the considerations of later parts of the paper favor the second interpretation, so that whatever reason one gives for the significant coefficient on model output, the fact remains that the real-business-cycle model driven by the Solow residual is not explanatory of output.

[10] We thank our colleague A. Colin Cameron for suggesting this explanation. He did so entirely on econometric grounds ignorant of the details of what the data were and how they were generated.

Table 3. Correlations of *YL*, *YM*, *SRL*

	YL	*YM*	*SRL*
YL	1.00	0.81	0.86
YM	0.81	1.00	0.98
SRL	0.86	0.98	1.00

In the next three sections, we provide evidence in support of our conjecture that the model output is not a genuine prediction of output but is in fact just output itself with an overlay of noise added by the filtering procedures and the model.

5. The Real Business Cycle Model Has Little or No Incremental Predictive Power

We conjecture that the real-business-cycle model itself provides little or no incremental information. Figure 2 demonstrates the plausibility of this claim. It shows that *YM* tracks *SRL* fairly closely, and that both track *YL* somewhat less closely. This point is quantified in Table 3, which shows that *YM* and *SRL* are highly correlated (correlation coefficient of 0.98), while each is substantially less, but still well, correlated with *YL*.

If our conjecture is correct, then both the construction of the Solow residual and the processing of it through the model would be sources of noise. The Solow residual itself (*SRL*) should then be a less noisy measure of *YL* and, not only should show the same ability to improve the fit of the univariate time-series model and induce serial correlation, but should dominate *YM* in the sense that it should carry all the information contained in *YM* and more and, therefore, should render *YM* insignificant in any regression in which it is also a regressor. Line 3 of Table 2 presents the results of a regression in which *SRL* replaces *YM* as an additional regressor. The *t*-statistic on the coefficient on *SRL* is 6.26—substantially larger than the *t*-statistic for *YM* (4.68) and the standard error of regression is lower than reported for the regression with *YM* in line 2: it shows a 14% reduction, where line 2 shows a 10% reduction, compared with line 1. This provides some indication in favor of our conjecture.

The matter is, however, somewhat more complicated. Line 4 reports a regression with both *YM* and *SRL* included as regressors. In this case, both variables are statistically significant and there is a further 7% reduction of the standard error of regression relative to line 1. To some extent the interpretation of line 3 is complicated by the multicollinearity of *YM* and *SRL*. This is not a complete explanation, however, and one might be justified in concluding that the model has incremental explanatory power after all, were it not for other evidence developed below.

Note that in all the regressions in Table 2, except the univariate time-series model in line 1, there is evidence of substantial serial correlation in the residuals.

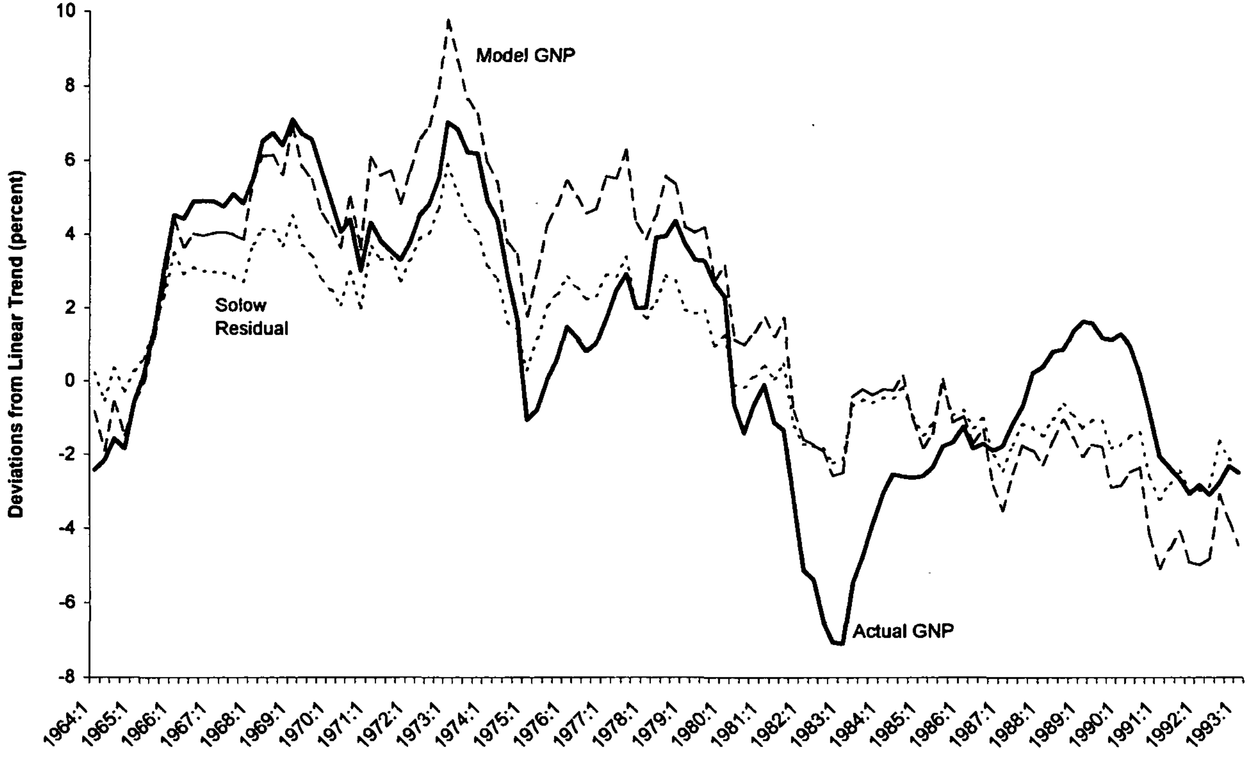

Fig. 2. Plot of actual GNP (*YL*), model GNP (*YM*), Solow residual (*SRL*); all series are linearly detrended.

6. The Solow Residual is Just Colored Noise...

The Solow residual is constructed by exploiting the supposed structural connection between factor inputs and output implied by a Cobb–Douglas production function. The idea is that the technology shocks that are used to drive the model are contained in the output series and can be extracted if careful account is taken of the behavior of inputs given technology. In this section we construct a *faux* (or fake) Solow residual by subtracting from the output a series that has similar stochastic properties to the factor input series but which bears absolutely no structural relationship to output. This series therefore imparts to the output series stochastic properties, similar to those of the Solow residual but, by construction, it is unable to isolate genuine technology shocks. We then use this *faux* Solow residual to generate model output and demonstrate that it behaves in precisely the same ways as the output of a model driven by the actual Solow residual.

Our conclusion is twofold: First, since the relationships of the driving process, the model output, and the actual data are the same, whether the driving process is the actual or the *faux* Solow residual, we have no reason to believe that the actual Solow residual carries information that the *faux* series clearly does not (i.e. it supports the view that the actual series is itself simply output plus colored noise, just like the *faux* series). Second, the value-added of the model is supposed to be that it converts the input of technology shocks into a good prediction of actual output. But the similarity in the predictive success of using actual or *faux* Solow residuals suggests that the success of the model cannot be attributed to an ability to process technology shocks; rather, it depends on the stochastic properties of the input series and the manner in which they are filtered by the model. This undermines any interpretation of Table 2, line 4, suggesting that the joint significance of the model output and the Solow residual are evidence of incremental predictive power on the part of the model.

To construct the *faux* series, which we call *FSRL* (the names of all *faux* series are just the names of the actual series with an F prefix), we first generated a series for total factor payments defined as:

$$PAY_t = \alpha K_t + (1-\alpha)L_t \tag{4}$$

where K_t is capital, L_t,is labor, and α is capital's share in output. Then we used this series to estimate the following time series model for *PAY*:

$$PAY_t = -3.31 + 0.0016TIME + v_t$$

$$\text{where } v_t = 1.24v_{t-1} - 0.29v_{t-2} + \eta_t. \qquad \eta_t \sim N(0,0.005) \tag{5}$$

The coefficients of this time series model were then used to generate a *faux* series, called *FPAY* in which the values of η_t were drawn from a random-number generator. Next, the constructed factor payment series was used to generate a Solow residual series; *FSR*. *FSR* was then linearly detrended and used as an input to the model exactly as described in Section 2.

Table 2, lines 5–7 present the regressions analogous to those reported in lines 2–4. The character of the results is identical. *Faux* model output is a significant predictor of actual output: it carries a high *t*-statistic and reduces the standard error of regression compared to line 1 by 25%; it induces substantial serial correlation in the residuals. The *faux* Solow residual is an even more significant predictor of actual output. In line 7, both *FYM* and *FSRL* are significant. In every case the relative orders of magnitude of coefficients and their *t*-statistics and the pattern of their signs are similar between the actual and the *faux* regressors.

7. ...and the Colours are Artificial

In this section, we provide further evidence that the apparent explanatory efficacy of the real-business-cycle model is illusory. In the last section, the *faux* Solow residual was constructed using an artificial series for factor payments. But since actual output was used, it was possible to compare the model output with the actual output. In this section, we create a completely simulated Solow residual in two stages. First, we estimated the following ARIMA(2,1,2) model for actual output:

$$\delta Y_t = 0.0026 - 0.156\Delta Y_{t-1} + 0.392\Delta Y_{t-2} + \omega_t$$

where $\omega_t = v_t + 0.49v_{t-1} - 0.05v_{t-2}$ and $v_t \sim N(0.000057)$.

Using its parameters and shocks drawn from a random number generator and accumulating from the actual level of *Y* in the fourth quarter of 1960, we created a simulated output series *SY* (the prefix *S* indicates simulated). *SY* was linearly detrended to generate *SYL*. A new realization of *FPAY* was generated as in Section 5 above. The simulated Solow residual was then constructed as $SSR = SY - FPAY$, and linearly detrended to yield *SSRL*. This series was then used as an input to the real-business-cycle model as in previous exercises.

Table 4 presents a set of regressions analogous to those in Table 2, lines 1–4. The key features of these regressions are the same as they are for the actual and the *faux* data in Table 2. Line 1 shows a univariate time-series model that fits well and shows no evidence of serial correlation. Line 2 shows that model output is statistically significant even after conditioning on the past history of *SYL*; and, just like the analogous case in Table 2, an improvement in the goodness of fit (9% as measured by the standard error of regression) is accompanied by evidence of significant serial correlation in the residuals. Line 3 shows that the simulated Solow residual, *SSRL*, is substantially more significant than *SYM* and improves the fit by well over double (26%). Yet, in line 4, when both *SYM* and *SRL* are included in the regressions, both are significant and both show the same sign pattern and relative magnitudes as the analogous regressions in lines 4 and 7 of Table 2.[11]

It is important to understand exactly what we have done here. It is not the case that we have created an artificial world in which there are artificial analogues to true technology shocks. For that to be true, *SY* and *FPAY* would have had to be connected by a production function. But they are not connected;

[11] The exact values reported in Table 4 are, of course, partly artifacts of the particular draw used to construct *SY* and *FPAY*. We have run many simulations, however, and the main features of the simulations reported here are robust.

Table 4. Regression results, simulated output

	Coefficient estimates[1]		Summary statistics			Regression diagnostics					
							Serial correlation up to order:[5]				
Line	Additional variables	Coefficient (t-statistics)	$\bar{R}^2$	SER^2	F statistic[3]	Normality[4]	1	2	4	Arch[6]	Hetero-skedasticity[7]
1	None	0.08 (3.22)	0.99	0.0079	1224.4	2.16	0.13	0.07	0.84	3.04	1.47
							(0.71)	(0.93)	(0.50)	(0.02)	(0.12)
							0.14	0.14	3.36	11.45	22.32
							(0.70)	(0.93)	(0.54)	(0.02)	(0.13)
2	*SYM*	0.18 (5.32)	0.99	0.0071	1366.9	1.60	21.54	11.53	7.60	1.26	1.00
							(0.00)	(0.00)	(0.00)	(0.29)	(0.47)
							19.76	21.06	26.67	5.05	18.12
							(0.00)	(0.00)	(0.00)	(0.28)	(0.45)
3	*SSRL*	0.45 (9.04)	0.99	0.0060	1910.2	0.77	8.56	12.58	6.24	0.67	1.02
							(0.00)	(0.00)	(0.00)	(0.62)	(0.45)
							8.74	22.61	22.82	2.73	18.40
							(0.00)	(0.00)	(0.00)	(0.60)	(0.43)

4	*SYM*	− 0.12 (2.36)	0.99	0.0059	1793.46	1.34	7.04	12.89	6.48	0.75	0.77
							(0.00)	(0.00)	(0.00)	(0.55)	(0.74)
	SSRL	0.63 (7.04)					7.35	23.24	23.71	3.08	16.26
							(0.00)	(0.00)	(0.00)	(0.54)	(0.70)

Each line reports a regression of *SYL* on a constant and eight own lags and additional variables using quarterly data, 1964:1–1993:1.

[1] Coefficient estimates are not reported for lags of *SYL* or the constant.

[2] Standard error of regression.

[3] Tests the null hypothesis that coefficients of all regressors are zero.

[4] Jarque–Bera test for normality of residuals.

[5] Breusch–Godfrey tests for serial correlation. The top pair of numbers is the F-form. The bottom pair is (number of observations) × R^2 form which is distributed asymptotically as Chi-square. *P*-values are given in parentheses.

[6] Tests the null of no autoregressive conditional heteroskedasticity. The top pair of numbers is the F-form. The bottom pair is the (number of observations) × R^2 form which is distributed asymptotically as Chi-square. *P*-values are given in parentheses.

[7] Tests null of no heteroskedasticity. The top pair of numbers is the F-form. The bottom pair is the (number of observations) × R^2 form which is distributed asymptotically as Chi-square. *P*-values are given in parentheses.

they are two independent series that have been created to possess time-series characteristics similar to their actual counterparts. The difference, *SY–FPAY*, cannot uncover technology shocks because the two series are unrelated and because there are no such shocks in *SY* to recover. Nevertheless, when *SSRL* is used as the driving process for the real-business-cycle model it shows that *SYM* carries statistically significant information about *SYL*. This must, by construction, be nothing more than an artifact of the filtering involved in detrending and in running *SSRL* through the model and of the stochastic properties of the root series. The relationship of actual output and model output cannot indicate that the model has captured a deep economic relationship; for there is no such relationship to capture. Rather, it shows that we are seeing a complicated version of a regression fallacy: output is regressed on a noisy version of itself, so it is no wonder that a significant relationship is found. That the model can establish such a relationship on simulated data demonstrates that it can do so with any data that are similar in the relevant dimensions. That it has done so for actual data hardly seems subject to further doubt.

8. Further Artifacts of Filtering: the H–P filter

We have shown so far that the close relationship of the model output to the actual data for output, results not from the model having high explanatory power but from the input to the model being nothing but a noisy measure of output itself and from the various filtering procedures that induce spurious correlations. This is clear enough in the case of linearly detrended data. The effect is even more pronounced when using data that has been passed through the Hodrick–Prescott (H–P) filter. The H–P filter is the most common detrending procedure used in the real-business-cycle literature.[12] If we follow the frequent practice of real-business-cycle modellers and run *YM* through the H–P filter to generate *YMHP* and then compare it with H–P-filtered *Y*, denoted *YHP*, the initial puzzle that motivated this paper is even more dramatic.

Table 5, line 1 presents a regression of *YHP* on a constant and 12 own lags.[13] Line 2 shows that when *YMHP* is entered as an additional regressor it comes in with a *t*-statistic of 12.56 and lowers the standard error of regression by 37%—much larger effects than reported for the comparison of *YM* to *YL*. (Just as in the case of the linearly detrended data, the additional regressors induce serial correlation in the residuals.) The root data in Tables 2 and 5 are the same; the reported results differ only in the filtering methods used to prepare the data for comparison.

[12] Hodrick & Prescott (1997) Let $x_t = \bar{x} + \hat{x}_t$ where $\bar{x}_t$, denotes the trend component and $\hat{x}_t$ the deviation from trend. Then the H–P filter chooses this decomposition in order to solve the following problem:

$$\min\left\{(1/T)\sum_{t=1}^{T}\hat{x}_t^2 + (\lambda/T)\sum_{t=2}^{T-1}[(\bar{x}_{t+1} - \bar{x}_t) - (\bar{x}_t - \bar{x}_{t-1})]^2\right\}$$

Following the universal practice of real-business-cycle modelers, we set $\lambda = 1600$ for quarterly data.

[13] We initially estimated a regression with only eight lags but moved to 12 to eliminate evidence of fourth-order serial correlation.

Table 5. Regression results, HP filtered output

	Coefficient estimates[1]		Summary statistics			Regression diagnostics					
							Serial correlation up to order:[5]				
Line	Additional variables	Coefficient (t-statistics)	$\bar{R}^2$	SER^2	F statistic[3]	Normality[4]	1	2	4	Arch[6]	Hetero-skedasticity[7]
1	None		0.81	0.0065	41.62	5.13	3.22	1.61	1.10	0.92	2.06
							(0.08)	(0.21)	(0.36)	(0.45)	(0.008)
							3.54	3.57	4.94	3.73	40.85
							(0.06)	(0.17)	(0.29)	(0.44)	(0.017)
2	*YMHP*	0.65 (12.56)	0.92	0.0041	108.50	5.23	9.82	6.00	3.92	0.30	0.97
							(0.002)	(0.003)	(0.005)	(0.88)	(0.52)
							10.27	12.42	15.98	1.23	25.56
							(0.001)	(0.002)	(0.003)	(0.87)	(0.49)
3	*SRHP*	0.95 (12.29)	0.92	0.0042	105.48	5.37	8.98	5.68	3.61	0.31	1.00
							(0.003)	(0.005)	(0.009)	(0.87)	(0.48)
							9.46	11.84	14.89	1.29	26.24
							(0.002)	(0.003)	(0.005)	(0.86)	(0.45)
4	*YMHP*	3.87 (3.06)	0.93	0.0040	106.59	4.13	27.19	13.59	13.29	0.36	0.83
							(0.000001)	(0.000006)	(0.00)	(0.84)	(0.71)
	SRHP	− 4.72 (2.55)					24.81	25.00	41.14	1.47	24.45
							(0.000001)	(0.000004)	(0.00)	(0.83)	(0.66)
5	*YM*	0.03 (1.84)	0.81	0.0065	39.56	7.46	12.14	6.08	3.03	1.12	2.11
							(0.0007)	(0.003)	(0.02)	(0.35)	(0.005)
							12.45	15.58	12.76	4.50	44.28
							(0.0004)	(0.002)	(0.01)	(0.34)	(0.014)

Table 5. *Continued*

	Coefficient estimates[1]		Summary statistics			Regression diagnostics					
							Serial correlation up to order:[5]				
Line	Additional Variables	Coefficient (t-statistics)	$\bar{R}^2$	SER^2	F statistic[3]	Normality[4]	1	2	4	Arch[6]	Hetero-skedasticity[7]
6	*SRL*	0.06 (2.20)	0.81	0.0064	40.21	7.53	15.61	7.79	3.95	1.07	1.97
							(0.0001)	(0.0007)	(0.005)	(0.38)	(0.00)
							15.53	15.63	16.10	4.30	42.48
							(0.00008)	(0.0004)	(0.003)	(0.37)	(0.022)
7	*YM*	−0.14 (1.64)	0.82	0.0064	38.14	6.49	18.61	9.21	4.60	0.73	2.08
							(0.00004)	(0.0002)	(0.003)	(0.37)	(0.022)
	SRL	0.29 (2.03)					18.20	18.20	18.50	2.99	46.61
							(0.00002)	(0.0001)	(0.001)	(0.56)	(0.01)

Each line reports a regression of *YHP* on a constant and twelve own lags and additional variables using quarterly data, 1964:1–1993:1.
[1] Coefficient estimates are not reported for lags of *YHP* or the constant.
[2] Standard error of regression.
[3] Tests the null hypothesis that coefficients of all regressors are zero.
[4] Jarque–Bera test for normality of residuals.
[5] Breusch–Godfrey tests for serial correlation. The top pair of numbers is the F-form. The bottom pair is (number of observations) × R^2 form which is distributed asymptotically as Chi-square. P-values are given in parentheses.
[6] Tests the null of no autoregressive conditional heteroskedasticity. The top pair of numbers is the F-form. The bottom pair is the (number of observations) × R^2 form which is distributed asymptotically as Chi-square. P-values are given in parentheses.
[7] Tests null of no heteroskedasticity. The top pair of numbers is the F-form. The bottom pair is the (number of observations) × R^2 form which is distributed asymptotically as Chi-square. P-values are given in parentheses.

Table 6. Correlations of *YHP, YMPH, SRHP*

	YL	*YM*	*SRL*
YHP	1.00	0.80	0.77
YMHF	0.80	1.00	0.99
SRHP	0.77	0.99	1.00

Table 5 shows that the patterns evident in Table 2 are reproduced when the data are H–P-filtered. Line 3 shows that the Solow residual itself has substantial explanatory power for *YHP*. When both *SRP* and *YMHP* are entered, both remain significant. One difference, however, is that *YMHP* carries a slightly higher *t*-statistic than *SRHP* throughout Table 5. This is, we believe for the reasons stated below, an artifact of the H–P filter.

First, although we have followed common practice in passing model output through the H–P filter, we are not aware of any published argument justifying this practice. Passing the actual data through the H–P filter is justified on the logic of the neoclassical growth model. In a steady state, the values of all stocks should have a common trend. Real-business cycle models take the trend itself as exogenous and aim to explain deviations from it. If the steady-state rate of growth is not constant over time, a slowly varying filter, like the H–P filter or a moving-average process, might capture the main movements in the steady state. By construction, the data generated by the model are stationary; there is no trend to extract.

Second, although we H–P-filtered the Solow residual for use as a regressor in Table 5, in order to better compare like with like, the model output was not generated using the H–P-filtered Solow residual as the measure of the technology shock (i.e. *SRL* is still the input to the model). This probably accounts for the fact that the *t*-statistics on *SRHP* are slightly lower than those on *YMHP* in Table 5.

The importance of the H–P filter is clearly visible when we compare the correlations among *YHP*, *YMHP* and *SRHP* in Table 6 with the correlations among *YHP*, *YM* and *SRL* in Table 7. The Solow residual and model output are closely correlated (correlation coefficients of 0.98 and 0.99) however they are filtered. The correlation coefficient between *SRL* and *YHP* is only 0.34, while the correlation coefficient between *SRHP* and *YHP* rises to 0.77. Similarly, the correlation coefficient between *YM* and *YHP* is only 0.30, while that between *YMHP* and *YHP* rises to 0.80.

Table 7. Correlations of *YHP*, *YM*, *SRL*

	YL	*YM*	*SRL*
YHP	1.00	0.30	0.34
YM	0.30	1.00	0.98
SRL	0.34	0.98	1.00

Essentially, the same message can be seen in Table 5, lines 5–7 in which non-H–P-filtered regressors are added to the univariate time-series model for *YHP*. Line 5 shows that *YM* is significant at the 10% but not the 5% level. Line 6 shows that *SRL* is more significant than *YM* (a *t*-statistic of 2.20 versus 1.84). Line 7 enters both *YM* and *YHP* as regressors. In this case *SRL* clearly dominates *YM*, although the regression has only a marginal explanatory advantage over the univariate model in line 1 (the standard error of regression drops trivially from 0.65 to 0.64).

Much of the explanatory power of the model data for the actual data is an artifact of the H–P filter, which massively raises the correlation between series that were not that closely related before filtering. Once the filtering of the model data is eliminated, model output has only a limited explanatory power for H–P filtered actual output, and the Solow residual completely dominates the model, suggesting that the independent contribution of the model is virtually nil. Still, using any combination of *YM* or *SRL* as regressors induces serial correlation.

Two points deserve notice. First, the general character of the results is consistent across Tables 2 and 5, which points to robust support for our conclusion that the real-business-cycle is not empirically successful in modelling output. Second, recalling that the root data is identical between Tables 2 and 5, much of the dramatic performance of *YMHP* in line 2, as well as some other features of the data, can be clearly attributed to adventitious characteristics of the H–P filter: it induces closer statistical relationships among data than the underlying economic relationships justify.[14]

9. Caveats and Conclusions

When Solow originally proposed the Solow residual, his purpose was to give a quantitative indication of the importance of long-run technological advancement.

[14] Kydland & Prescott (1995, pp. 9–10) defend the use of the H–P filter against critics who have argued that it induces spurious cycles by stating that deviations from trends defined by the H–P filter 'measure nothing' but instead are 'nothing more than well-defined statistics'; and, since 'business cycle theory treats growth and cycles as being integrated, not as a sum of two components driven by different factors', 'talking about the resulting statistics as imposing spurious cycles makes no sense'. The logic of Kydland & Prescott's position escapes us. It is true that real-business-cycle theory treats the business cycle as the equilibrium adjustments of a neoclassical growth model subject to technology shocks. In practice, real-business-cycle models are calibrated on the assumption that steady-state values for key ratios should conform to their sample averages. The models are linearized around the steady state so that the output of the model is expressed as deviations from the steady state. Generally, if actual values are compared with the model output, the actual values are first detrended using an ad hoc filter, which imposes the empirical assumption that the filter's trend is a good approximation of the true steady state. If instead, comparisons are made of the levels of actual and modelled variables, this same ad hoc trend is added to the deviations from steady state to generate modelled levels. In either case, the relevant steady-state is not jointly modelled with the deviations from steady-state, but is generated from the ad hoc filter. That this is the practice does not say that such a joint modelling exercise could not be done in theory. That it is not actually done in practice means that the objection to the H–P filter raised by many critics remains cogent. Our work, and that of the critics that Kydland & Prescott wish to dismiss, demonstrates that the choice of which ad hoc method is used to extract the balanced-growth path greatly affects the stochastic properties of the modelled variables and their relationships with the actual data.

It may well be useful in this role. When Prescott (1985) proposed to use the Solow residual, his purpose was to capture short-run technology shocks. Let us now draw the threads of our investigation together to show why we have reason to doubt that it can be successfully used for that purpose.

The starting point of our investigation is a two-part empirical puzzle. First, when the model data for output are added as a regressor to a profligately parameterized univariate time-series model for actual output, model output comes in highly significantly and improves the fit of the regression markedly. Second, model output induces substantial serial correlation in the data. We want to discriminate between two explanations for these phenomena. On the one hand, the fit might mean that the model is so good that the model output carries substantial non-redundant information about actual output. On the other hand, it could be that model output is nothing more than a noisy measure of actual output, and the regression amounts to regressing output on itself. The induced serial correlation favors that explanation.

For a model to be explanatory it must add information beyond that which is contained in the exogenous processes that it takes as inputs. Since the only input to the typical real-business-cycle model is the Solow residual, it is easy to check whether the model adds information: we simply use the Solow residual itself as a regressor. We discovered that the Solow residual shows the same pattern of improving the fit and inducing serial correlation.

The only fly in the ointment is that the t-statistics remain significant on both the Solow residual and the model output when both are entered as regressors, perhaps suggesting that it has incremental explanatory power relative to the Solow residual. This result should be discounted. First, t-statistics are biased upward in the presence of serially correlated errors. Second, that this result is a statistical artifact is reinforced by the regressions using a *faux* Solow residual as the driving process. The model is supposed to have an explanatory advantage because it uses calibrated economic theory to process information about technology shocks. The *faux* Solow residual does not carry such information by its very construction, yet all the same interrelationships appear between actual output and the *faux* Solow residual and *faux* model output. In fact, all the same relationships appear even when entirely simulated output and factor input series are used that have no economic relationship to each other.

The results for the *faux* series also underwrite the case for the Solow residual simply not carrying any relevant information on technology shocks. The *faux* Solow residual mimics the time-series properties of the actual Solow residual, but by its construction does not carry any information about technology shocks. Nevertheless, the performance of it and the other *faux* variables that are created by using it as an input to the model or to the H–P filter is qualitatively the same as that using the actual Solow residual. The *faux* Solow residual is, by definition, nothing but colored noise. The actual Solow residual is thus indistinguishable from colored noise.

It is important to understand the limited nature of our criticism of the real-business-cycle model. We do not claim to have demonstrated that the real-business-cycle model is intrinsically inadequate. Our evidence is consistent with the critical work of other investigators. Cogley & Nason (1995b) show that in models like the real-business-cycle model of this paper the time-series properties of model output—e.g. the autocorrelation properties that are usually used to evaluate the models—are inherited from the time-series properties of the Solow residual and not contributed by the model itself. Similarly, Watson's (1993) spectral analysis shows that typical real-business-cycle models lack power at business cycle frequencies. These results are consistent with our findings, but we do not address them directly. It may be that with a better measure of technology shocks the real-business-cycle model would have substantial explanatory power. It may also be that real-business-cycle models are inadequate as explanations of output, but nevertheless provide explanations of consumption, investment, and hours worked conditional on output. We do not rule that out.

What we must finally conclude is that the construction of the Solow residual and the manner in which it is processed by the model and by the linear and H–P filters ensure that the apparent explanatory power of the real-business-cycle model for actual historical business cycles is a pure artifact: the model output is nothing more than actual output itself with an overlay of complicated statistical noise. Real-business-cycle models in practice provide no explanation whatsoever of the course of actual business-cycle history.

References

Benhabib, J., Rogerson, R. & Wright, R. (1991) Homework in macroeconomics: household production and aggregate fluctuations, *Journal of Political Economy* 6, pp. 1166–1187.

Canova, F. (1991) Detrending and business cycle facts, Unpublished typescript, Department of Economics, European University Institute, Florence.

Canova, F., Finn, M. & Pagan, A.R. (1994) Evaluating a real business cycle model, in: C. Hargreaves (Ed.) *Nonstationary Time Series Analysis and Cointegration* (Oxford, Oxford University Press).

Christiano, L.J. & Eichenbaum, M. (1992) Current real-business cycle theories and aggregate labor market fluctuations, *American Economic Review* 82, pp. 430–450.

CITIBASE: Citibank Economic Database (New York, Citibank, N.A.) 1978, July 1993 update.

Cogley, T. & Nason, J.M. (1995a) Effects of the Hodrick–Prescott filter on trend and difference stationary time series: implications for business cycle research, *Journal of Economic Dynamics and Control*, 19, pp. 253–278.

Cogley, T. & Nason, J.M. (1995b) Output dynamics in real business cycle models, *American Economic Review*, 85(3), pp. 492–511.

Cooley, T.F. (1995) Contribution to a conference panel discussion: what do we know about how monetary policy affects the economy? At the *19th Annual Economic Policy Conference*, Federal Reserve Bank of St. Louis, *Federal Reserve Bank of St. Louis Review*, 77(3), pp. 131–137.

Cooley, T.F. and Prescott, E.C. (1994) Econornic growth and business cycles, in: T.F. Cooley (Ed.) *Frontiers of Business Cycle Research* (Princeton, Princeton University Press).

Fair, R.C. & Schiller, R.J. (1990) Comparing information in forecasts from econometric models, *American Economic Review*, 80(3), pp. 375–389.

Farmer, R.E.A. (1993) *The Macroeconomics of Self-fulfilling Prophecies* (Cambridge, MA, MIT Press).

Granger, C.W.J & Newbold, P. (1986) *Forecasting Economic Time Series*, 2nd edn. (New York, Academic Press).

Hall, R.E. (1986) Market structure and macroeconomics fluctuations, *Brookings Papers on Economic Activity*, 86, pp. 265–338.

Hall, R.E. (1990) Invariance Properties of Solow's productivity residual, in: P. Diamond (Ed.)

Growth/Productivity/Unemployment: essays to celebrate Bob Solow's birthday, pp. 71–112 (Cambridge, MIT Press).
Hansen, G.D. (1985) Indivisible labor and the business cycle, *Journal of Monetary Economics*, 16, pp. 39–69.
Hansen, G.D. & Prescott, E.C. (1993) Did technology shocks cause the last recession, *American Economic Review*, 83, pp. 280–286.
Hansen, L.P. & Heckman, J.J. (1996) The empirical foundations of calibration, *Journal of Economic Perspectives*, 10, pp. 87–104.
Hartley, J.E. (1994a) Technology in macroeconomic models, Ph.D. Dissertation, Department of Economics, University of California, Davis, June 1994.
Hartley, J.E. (1994b) Does the Solow residual actually measure changes in technology, unpublished typescript, Mt. Holyoke College, October.
Hodrick, R.J. & Prescott, E.C. (1997) Post-war U.S. business cycles: an empirical investigation, *Journal of Money, Credit, and Banking*, 29, pp. 1–16.
Hoover, K.D. (1994) Six queries about idealization in an empirical context, *Ponzan Studies in the Philosophy of the Sciences and the Humanities*, 38, pp. 43–53.
Hoover, K.D. (1995) Facts and artifacts: calibration and the empirical assessment of real-business-cycle models, *Oxford Economic Papers* 47, pp. 24–44.
King, R.G., Plosser, C.I. & Rebelo, S.T. (1988) Production, growth, and business cycles I: the basic neoclassical model, *Journal of Monetary Economics*, 21, pp. 309–42.
Kydland, F.E. & Prescott, E.C. (1982) Time to build and aggregate fluctuations, *Econometrica*, 50, pp. 1345–1370.
Kydland, F.E. & Prescott, E.C. (1991) the econometrics of the general equilibrium approach to business cycles, *Scandinavian Journal of Economics*, 93(2), pp. 161–178.
Kydland, F.E. & Prescott, E.C. (1995) The computational experiment: an econometric tool, Federal Reserve Bank of Minneapolis, Research Department, Staff Report No 178.
Lucas, R.E. (1987) *Models of Business Cycles* (Oxford, Blackwell).
Lucas, R.E., Jr. (1981) Methods and problems in business cycle theory, in: *Studies in Business-Cycle Theory*, pp. 271–296 (Oxford, Blackwell).
McCallum, B.T. (1989) Real business cycle models, in: R.J. Barro (Ed.) *Modern Business Cycle Theory* (Oxford, Blackwell).
Musgrave, J.C. (1992) Fixed reproducible tangible wealth in the United States: revised estimates, *Survey of Current Business*, 72, pp. 106–107.
Plosser, C.I. (1989) Understanding real business cycles, *Journal of Economic Perspectives*, 3, pp. 51–78.
Prescott, E.C. (1983) Can the cycle be reconciled with a consistent theory of expectations? or a progress report on business cycle theory, Working Paper No. 239, Research Department, Federal Reserve Bank of Minneapolis.
Prescott, E.C. (1986) Theory ahead of business cycle measurement, *Federal Reserve Bank of Minneapolis Quarterly Review*, 10, pp. 9–22.
Solow, R.M. (1957) Technical change and the aggregate production function, *Review of Economics and Statistics*, 39, pp. 312–320.
Watson, M.W. (1993) Measures of fit for calibrated models, *Journal of Political Economy*, 101, pp. 1011–1041.

Appendix

Model Description

For our analysis, we use a real business cycle model with divisible labor which is identical (except for the particular parameter values used for the simulations) to that described in Hansen (1985). That is, we assume the following social planner problem:

$$\max \quad E_0\left[\sum_{t=0}^{\infty} \beta^t(\ln c_t + A \ln(l - h_t))\right]$$

$$\begin{aligned} s.t. \quad & c_t + i_t \leq z_t k_t^{\alpha} h_t^{l-\alpha} \\ & k_{t+l} = k_t(l-\delta) + i_t \\ & \ln z_{t+1} = \rho \ln z_t + \varepsilon_{t+1} \end{aligned} \tag{P1}$$

where E denotes the expectations operator, β agents' discount factor, c_t is consumption, h_t denotes labor, i_t is investment, k_t is the capital stock, z_t is the shock to technology, δ is the depreciation rate of capital, and ε is the innovation to the technology shock (assumed to be log normally distributed with mean 0 and standard deviation σ_ε.

In order to solve for the equilibrium laws of motion for the endogenous variables implied by the above problem, the numerical approximation procedure described in Farmer (1993) was used. This method involves taking first-order Taylor series expansions of the necessary conditions associated with (P1) around the steady-state values (defined by setting $z_t = 1$) of c, h and k. The resulting linear structure implies that the solution values for next-period capital, current consumption and labor (the control variables) will be linear functions of the current state variables (the beginning-of-period capital stock and the current shock to technology (Both control and state variables are expressed as percentage deviations from steady-state values). The parameters of these functions are determined by imposing the transversality condition associated with (P1). Once these functions are determined, the equilibrium behavior of output and investment are determined by the first-order approximations of the production function and the economy-wide resource constraint. The actual equations we used to generate the data are presented below in the discussion of the calibration of the model.

Data Construction

In order to parameterize the model described above, we followed the approach described in Cooley & Prescott (1994) which attempts to impose a high level of consistency between model constructs and measured data. For instance, in the model output, is a function of the aggregate capital stock; hence, for consistency with the data it is necessary to construct imputed income flows from the stock of consumer durables as well as that from government capital. Also, since the model does not contain inventories or a foreign sector, the measurement of aggregate output should reflect these assumptions. We briefly describe the measurement and construction of the data below; for a more detailed description, the reader is referred to Cooley & Prescott (1994).

Income from the private capital stock is defined as:

$$Y_{\mathrm{pk}} \; (r + \delta_{\mathrm{pk}})K_{\mathrm{p}} \tag{B1}$$

where (Y_{pk}, δ_{pk}, K_p, r) denote the income from fixed private capital, the depreciation rate of private capital, the stock of private capital, and the return on private capital, respectively. Given measures of the first three variables, an estimate of r can be computed using Equation (B1). This estimate, along with an estimate of the depreciation rates of government capital and consumer durables can then be used to compute the income flows from the stock of consumer durables and government capital.

First, however, income from private capital must be determined, which necessitates an estimate of the fraction of the proprietor's income due to the existing private capital stock. To do this, it is assumed that private capital's share of total GNP, θ_p, is the same as that for the proprietor's income (net of taxes and subsidies). That is,

$$Y_{\mathrm{pk}} = (RI + CP + NETINT) + \theta_{\mathrm{p}}(PI + NNP - NI) + (GNP - NNP) = \theta_{\mathrm{p}} GNP \tag{B2}$$

RI denotes rental income, *CP is corporate profits, NETINT* is net interest, *PI* is proprietor's income, *NNP* is net national product, *NI* is national income, and *GNP* is gross national product. Hence the first term in parentheses is unambiguous capital income, while the third term in parentheses measures depreciation. Using the right-hand side equality, θ_p can be expressed as a function of the various income flows. Once this share is calculated, multiplying by *GNP* determines Y_{pk}. Then, using this estimate in Equation (B1) determines r, the return on capital.

Next, time series for the depreciation rates of government capital and consumer durables were constructed using the law of motion for both capital stocks:

$$\delta_{\mathrm{gk},\,t}=\frac{(K_{\mathrm{g},\,t}+I_{\mathrm{g},\,t}-K_{\mathrm{g},\,t+1})}{K_{\mathrm{g},\,t}} \tag{B3}$$

$$\delta_{\mathrm{cd},\,t}=\frac{(K_{\mathrm{cd},\,t}+I_{\mathrm{cd},\,t}-K_{\mathrm{cd},\,t+1})}{K_{\mathrm{cd},\,t}} \tag{B4}$$

The averages of both time series for the depreciation rates were identified as the constant depreciation rates, δ_g and δ_{cd}. With measures of the stocks of government capital and consumer durables, the estimates of r, δ_g and δ_{cd} can be used to construct their imputed income flows via Equation (B1).

Then, total income, TY, is measured as:

$$TY=GNP+Y_{\mathrm{gk}}+Y_{\mathrm{cd}} \tag{B5}$$

As a final adjustment, net exports are included in investment expenditures since the model economy does not include a foreign sector. Hence, measured investment (the empirical counterpart to investment in the model) is:

$$I=FPI+\Delta INV+GI+ICD+NEX \tag{B6}$$

Where FPI is fixed private investment, ΔINV is the change in business inventories, GI denotes government investment in durables and structures, ICD is purchases of consumer durables, and NEX denotes net exports.

Data

The following data over the sample period 1960:1–1993:1 were used in constructing the above variables.

Government investment. Citibase series: *ggndq* (defense durables), *ggncq* (defense structures), *ggodq* (non-defense durables), *ggocq* (non-defense structures), *ggsdq* (state and local durables), *ggscq* (state and local structures). The sum of these series equals GI.

Determining θ_p private capital's share of income. The following Citibase series were used (all nominal denominated series were deflated by the implicit GNP price deflator denoted as *price*): *gprenj* = rental income, *gpjva* = corporate profits, *gnint* = net interest income, *gproj* = proprietor's income, *gmp* = net national product, *gy* = national income, *gnpq* = real gnp. Hence, θ_p is determined by:

$$\theta_{\mathrm{p}}=\frac{\dfrac{(gprenj+ghpva+gnint)}{\mathrm{price}}+\left(gnpq-\dfrac{gnnp}{price}\right)}{gnpq-\left(\dfrac{gproj+gnnp-\mathrm{gy}}{\mathrm{price}}\right)} \tag{C1}$$

Constructing the income flows from government capital and the stock of consumer durables. Measures for the net stock of reproducible private capital (residential and non-residential), the net stock of consumer durables, and the net stock of government capital are all taken from Musgrave (1992). These are annual series. For the stock of inventories, the annual average of the Citibase series *glq* was used. The total capital stock, TK is determined by

$$TK=K_{\mathrm{p}}+glq+K_{\mathrm{cd}}+K_{\mathrm{g}} \tag{C2}$$

Using the estimate of θ_p determined by Equation (C2), the average rate of return on private capital, r, was measured as 9.3% (expressed as an annual rate). Government investment was measured by GI as defined above while the purchases of consumer durables (Citibase series *gcdq*) was used for investment in consumer durables. Using the annual average of these series in conjunction with the respective measures of the relevant capital stock produced the estimates of θ_{cd} and θ_{gk} (see

Equations (B3) and (B4)). These were 16.7% and 4.5% respectively. Using these values in Equation (Bl) produced the imputed income flows from the stock of consumer durables, denoted Y_{cd}, and from government capital, denoted Y_{gk}. Total income was measured as:

$$TY = gnpq + Y_{\text{cd}} + Y_{\text{gk}} \tag{C3}$$

Then capital's share of total income, the parameter α in the production function, was measured as:

$$\alpha = \frac{Y_{\text{pk}} + Y_{\text{cd}} + Y_{\text{gk}}}{TY} \tag{C4}$$

The average of α was 0.41 over the sample period.

Investment. Along with the measure of government investment, *GI*, the following Citibase series were used: fixed private investment (*gifq*), change in business inventories (*gvq*), purchases of consumer durables (*gcdq*), and net exports (*netex*). The sum of these series was used to measure total investment

Labor. Labor input was measured by the Citibase series *lhours*.

Consumption. The sum of the Citibase series purchases of non-durables (*gcnq*) and services (*gcsq*) was used to measure consumption.

Finally, total output, total investment, labor, consumption, and the total capital stock were expressed in per-capita terms by dividing by the adult civilian non-institutional population (Citibase series *pm*20 + *pf*20). These series were then converted to natural logarithms.

Calibrating the Parameter Values

In order to solve the model described in Section 2, parameter values for agents' preferences (β, A) as well as technology (α, δ, ρ) must be specified. As mentioned above, α represents capital's share which for the sample period we studied was equal to 0.41. The agents' discount factor was set to 0.984 while the depreciation rate was assumed to be 0.018; these values imply that the steady-state capital–output ratio is equal to 12, roughly duplicating that in the data. The weight on utility from leisure (A) was set to 2.15—this implies that 26% of time is spent in work activities again matching the sample average.

The remaining parameter, ρ, was determined by first introducing deterministic technological growth into the above model. (While this implies that the solution variables in (P1) will be growing over time, a transformation of the variables which removes the secular trend permits the solution method described above to be applied to the transformed problem, see King *et al.*, 1988.) That is, the technology shock was defined as:

$$Z_t = \lambda^t z_t \tag{Dl}$$

Z_t is measured as the Solow residual:

$$Z_t = y_t - \alpha k_t - (1 - \alpha)h_t \tag{D2}$$

(All variables on the right-hand side are measured in logarithms.) Then, z_t was identified as the residuals from regressing Z_t on a linear time trend. Analysing the autoregressive properties of this series produced an estimate for ρ of 0.96.

Using these values produced the following equilibrium laws of motion:

$$\hat{c}_t = 0.634\, \hat{k}_t + 0.02\, \hat{h}_t + 0.38\, z_t$$

$$\hat{k}_{t+1} = -0.066\, \hat{c}_t + 1.016\, \hat{k}_t + 0.049\, \hat{h}_t + 0.084\, z_t$$

$$\hat{h}_t = 1.32\, \hat{c}_t + 0.54\, \hat{k}_t + 1.32\, z_t$$

$$\hat{y}_t = z_t + 0.41\,\hat{k}_t + 0.59\,\hat{h}_t$$

$$\hat{i}_t = \left(\frac{1.45}{0.31}\right)\hat{y}_t - \left(\frac{1.14}{0.31}\right)\hat{c}_t$$

The caret denotes percentage deviation; from the steady-state. The first three equations are the solutions generated by the linear-approximation method described above. The fourth equation is the linearized form of the production function while the last equation is the first-order approximation of the aggregate resource constraint. The numerator of the first term in parentheses is the steady-state value of output, the numerator of the second term is the steady-state value of consumption, while the term in the denominator is the steady-state value of investment.