

Reformulating Empirical Macro-econometric Modelling

David F. Hendry

Economics Department, Oxford University

and

Grayham E. Mizon*

Economics Department, Southampton University.

Abstract

The policy implications of estimated macro-econometric systems depend on the formulations of their equations, the methodology of empirical model selection and evaluation, the techniques of policy analysis, and their forecast performance. Drawing on recent results in the theory of forecasting, we question the role of 'rational expectations'; criticize a common approach to testing economic theories; show that impulse-response methods of evaluating policy are seriously flawed; and question the mechanistic derivation of forecasts from econometric systems. In their place, we propose that expectations should be treated as instrumental to agents' decisions; discuss a powerful new approach to the empirical modelling of econometric relationships; offer viable alternatives to studying policy implications; and note modifications to forecasting devices that can enhance their robustness to unanticipated structural breaks.

JEL classification: C3, C5, E17, E52, E6

Keywords: economic policy analysis, macro-econometric systems, empirical model selection and evaluation, forecasting, rational expectations, impulse-response analysis, structural breaks.

*Published in OXFORD REVIEW OF ECONOMIC POLICY, Volume 16, No.4, WINTER 2000, pages 138-159.

Financial support from the U.K. Economic and Social Research Council under grant L138251009 is gratefully acknowledged. We are pleased to acknowledge helpful comments from Chris Allsopp, Mike Clements, Jurgen Doornik, Bronwyn Hall, John Muellbauer and Bent Nielsen.

1 Introduction

The policy implications derived from any estimated macro-econometric system depend on the formulation of its equations, the methodology used for the empirical modelling and evaluation, the approach to policy analysis, and the forecast performance. Drawing on recent results in the theory of forecasting, we question the role of ‘rational expectations’ in the first stage; then criticize the present approach to testing economic theories prevalent in the profession; next, we show that impulse-response methods of evaluating the policy implications of models are seriously flawed; and finally, question the mechanistic derivation of forecasts from econometric systems. In their place, we propose that expectations should be treated as instrumental to agents’ decisions; suggest a powerful new approach to the empirical modelling of econometric relationships; offer viable alternatives to studying policy implications; and discuss modifications to forecasting devices that can enhance their robustness to unanticipated structural breaks. We first sketch the arguments underlying our critical appraisals, then briefly describe the constructive replacements, before presenting more detailed analyses of these four issues. Sub-section 1.1 summarizes our critiques, and sub-section 1.2 introduces our remedies.

1.1 Four critiques of present practice

Our approach builds on extensive research that has radically altered our understanding of the causes of forecast failure, the occurrence of which was one of the driving forces behind the so-called ‘rational expectations revolution’ that replaced ‘Keynesian’ models. Forecast failure is a significant deterioration in forecast performance relative to the anticipated outcome, usually based on the historical performance of a model: systematic failure is the occurrence of repeated mis-forecasting. The research reveals that the causes of forecast failure differ from what is usually believed – as do the implications. To explain such differences, we begin by reconsidering the ‘conventional’ view of economic forecasting.

When the data processes being modelled are weakly-stationary (so means and variances are constant over time), three important results can be established. First, causal variables will outperform non-causal (i.e., variables that do not determine the series being forecast), both in terms of fit and when forecasting. Secondly, a model that in-sample fully exploits the available information (called congruent) and is at least as good as the alternatives (encompassing) will also dominate in forecasting; and for large samples, will do so at all forecast horizons. Thirdly, forecast failure will rarely occur, since the sample under analysis is ‘representative’ of the sample that needs to be forecast – moreover that result remains true for mis-specified models, inaccurate data, inefficient estimation and so on, so long as the process remains stationary. Such theorems provide a firm basis for forecasting weakly-stationary time series using econometric models: unfortunately, they can be extended to non-stationary processes only when the model coincides with the data generation process (DGP).

The systematic mis-forecasting and forecast failure that has periodically blighted macroeconomics highlights a large discrepancy between such theory and empirical practice, which is also visible in other disciplines: see e.g., Fildes and Makridakis (1995) and Makridakis and Hibon (2000). The key problem is the inherently non-stationary nature of economic data – even after differencing and cointegration transforms have removed unit roots – interacting with the impossibility in a high-dimensional and evolving world of building an empirical model which coincides with the DGP at all points in time. Consequently, one can disprove the most basic theorem *that forecasts based on causal variables will dominate those from non-causal*. Restated, it is easy to construct examples where forecasts based on variables that do not enter the DGP outperform those based on well-specified causally-sound models – one such example is shown below. Importantly, such results match the empirical evidence: we have

opened Pandora's Box, with profound implications that are the focus of this paper.

Having allowed the data process to be non-stationary and models to be mis-specified representations thereof (both in unspecified ways), one might imagine that an almost indefinite list of problems could precipitate forecast failure. Fortunately, that is not the case. To understand why, we must dissect the ingredients of econometric models. In general, econometric models have three main components: deterministic terms, namely variables whose future values are known (such as intercepts which are 1, 1, 1,... and trends, which are 1, 2, 3, 4...); observed stochastic variables with known past, but unknown future, values (such as GNP and inflation); and unobserved errors all of whose values (past, present and future) are unknown. Most relationships in models involve all three components because that is how we conceive of the data. In principle, any or all of the components could be: mis-specified; poorly estimated; based on inaccurate data; selected by inappropriate methods; involve collinearities or non-parsimonious formulations; and suffer structural breaks. Moreover, forecast failure might result from each 'problem'. Given the complexity of modern economies, most of these 'problems' will be present in any empirical macro-model, and will reduce forecast performance by increasing inaccuracy and imprecision. However, and somewhat surprisingly, most combinations do not in fact induce systematic forecast failure.

The taxonomy of sources of forecast errors in Clements and Hendry (1998, 1999a) implicates unanticipated forecast-period shifts in deterministic factors (such as equilibrium means, examples of which are the means of the savings rate, velocity of circulation, and the NAIRU) as the dominant cause of systematic failure. As explained in section 2.1, there is an important distinction between shifts in the deterministic components (such as intercepts) that enter models, and those that precipitate forecast failure (unmodelled shifts in data means), but for the moment we leave that to one side, as the former is usually sufficient for the latter. The crucial converse is that forecast failure is not in fact primarily due to the list of 'problems' in the previous paragraph, or even the Lucas (1976) critique of changing parameters: by themselves, none of these factors induces systematic failure. Our first critique now follows – since 'rational expectations' claim to embody the actual conditional expectations, they do not have a sound theoretical basis in an economy subject to deterministic shifts. Further, in the presence of unmodelled deterministic shifts, models embodying previously-rational expectations will not forecast well in general.

Turning to the second critique, tests of economic theories based on whole-sample goodness of fit comparisons can be seriously misled by unmodelled deterministic shifts. This occurs because such shifts can be proxied by autoregressive dynamics, which has two implications. First, deterministic shifts induce apparent unit roots, so cointegration often fails in the face of such breaks. Thus, long-run relationships – often viewed as the statistical embodiment of economic theory predictions – then receive no support. Secondly, such false unit roots can make lagged information from other variables appear irrelevant, so tests of theories – particularly of Euler equations – can be badly distorted. Our second critique now follows: so long as the degree of non-congruence of a model is unknown, false theories can end being accepted, and useful ones rejected.

A necessary condition for both economic theories and macro-economic models to be of practical value is that their parameters remain constant over the relevant horizon, and for the admissible range of policy changes to be implemented. Many structural breaks are manifest empirically, and so are easy to detect; deterministic shifts are a salient example. The class of breaks that are easy to detect comprises shifts in the unconditional expectations of non-integrated (denoted $I(0)$) components. Their ease of detection is the obverse of their pernicious effect on forecast performance. However, it transpires that a range of parameter changes in econometric models cannot be easily detected by conventional statistical tests. This class includes changes that leave unaltered the unconditional expectations, even

when dynamics, adjustment speeds, and intercepts are radically altered: illustrations are provided in Hendry and Doornik (1997). This leads to our third critique – impulse-response methods of evaluating the policy implications of models are dependent on the absence of such ‘undetectable breaks’, and so can be mis-leading in both sign and magnitude when non-deterministic shifts have occurred, even when models are rigorously tested (and certainly when minimal testing occurs).

Fourthly, there is ample evidence that forecasts from econometric systems can err systematically in the face of deterministic shifts, such that they perform worse than ‘naive’ methods in forecasting competitions. Theory now exists to explain how and why that occurs: see e.g., Clements and Hendry (1999c). The implementation of cointegration may in practice have reduced the robustness of econometric-model forecasts to breaks, by ensuring they adjust back to pre-existing equilibria, even when those equilibria have shifted. Mechanistic econometric-model based forecasts, therefore, are unlikely to be robust to precisely the form of shift that is most detrimental to forecasting. It is well known that devices such as intercept corrections can improve forecast performance (see e.g., Turner, 1990), but manifestly do not alter policy implications; and conversely, that time-series models with no policy implications might provide the best available forecasts. Hence our fourth critique – it is inadvisable to select policy-analysis models by their forecast accuracy: see Hendry and Mizon (2000).

1.2 Some remedies

The existence of these four problems implies that many empirical macro-econometric models are incorrectly formulated and wrongly selected, with policy implications derived by inappropriate methods. Whilst we suspect that some amelioration arises in practice as a result of most macro-forecasters continuing to use intercept corrections to improve forecasts, the almost insuperable problems confronting some approaches to macro-economics remain. Fortunately though, effective alternatives exist.

First, since expectations are instrumental to the decisions of economic agents, not an end in themselves, the devices that win forecasting competitions – which are easy to use and economical in information – suggest themselves as natural ingredients in agents’ decision rules (possibly ‘economically-rational expectations’: see Feige and Pearce, 1976). We show that is the case, with the interesting implication that the resulting rules may not be susceptible to the Lucas (1976) critique, thus helping to explain its apparent empirical irrelevance: see Ericsson and Irons (1995).

Next, stimulated by Hoover and Perez (1999), Hendry and Krolzig (1999) investigate econometric model selection from a computer-automation perspective, focusing on general-to-specific reduction approaches, embodied in the program *PcGets* (general-to-specific: see Krolzig and Hendry, 2000). In Monte Carlo experiments, *PcGets* recovers the DGP with remarkable accuracy, having empirical size and power close to what one would expect if the DGP were known, suggesting that search costs are low. Thus, a general-to-specific modelling strategy that starts from a congruent general model and requires congruence and encompassing throughout the reduction process offers a powerful method for selecting models. This outcome contrasts with beliefs in economics about the dangers of ‘data mining’. Rather, it transpires that the difficult problem is not to eliminate spurious variables, but to retain relevant ones. The existence of *PcGets* not only allows the advantages of congruent modelling to be established, it greatly improves the efficiency of modellers who take advantage of this model-selection strategy.

Thirdly, there is a strong case for using open, rather than closed, macro-econometric systems, particularly those which condition on policy instruments. Modelling open systems has the advantage that amongst their parameters are the dynamic multipliers which are important ingredients for estimating the responses of targets to policy changes. Further, it is difficult to build models of variables such as interest rates, tax rates, and exchange rates, which are either policy instruments or central to the determination

of targets. Since many policy decisions entail shifts in the unconditional means of policy instruments, corresponding shifts in the targets' unconditional means are required for policy to be effective. The relevant concept is called co-breaking, and entails that although each variable in a set shifts, there are linear combinations that do not shift (i.e., are independent of the breaks: see Clements and Hendry, 1999a, ch. 9). Co-breaking is analogous to cointegration where a linear combination of variables is stationary although individually they are all non-stationary. Whenever there is co-breaking between the instrument and target means, reliable estimates of the policy responses can be obtained from the model of the targets conditioned on the instruments, despite the probable absence of weak exogeneity of policy instruments for the parameters of interest in macro-models (due to mutual dependence on previous disequilibria): Ericsson (1992) provides an excellent exposition of weak exogeneity. The existence of co-breaking between the means of the policy instruments and targets is testable, and moreover, is anyway necessary to justify impulse-response analysis (see Hendry and Mizon, 1998).

Finally, there are gains from separating policy models – to be judged by their ability to deliver accurate advice on the responses likely from policy changes – from forecasting models, to be judged by their forecast accuracy and precision. No forecast can be robust to unanticipated events that occur after its announcement, but some are much more robust than others to unmodelled breaks that occurred in the recent past. Since regime shifts and major policy changes act as breaks to models that do not embody the relevant policy responses, we discuss pooling robust forecasts with scenario differences from policy models to avoid both traps.

We conclude that the popular methodologies of model formulation, modelling and testing, policy evaluation, and forecasting may prejudice the accuracy of implications derived from macro-econometric models. Related dangers confronted earlier generations of macro-models: for example, the use of dynamic simulation to select systems was shown in Hendry and Richard (1982) to have biased the choice of models to ones which over-emphasized the role of unmodelled ('exogenous') variables at the expense of endogenous dynamics, with a consequential deterioration in forecast performance, and mis-leading estimates of speeds of policy responses. As in that debate, we propose positive antidotes to each of the major lacunae in existing approaches. The detailed analyses have been presented in other publications: here we seek to integrate and explain their implications for macro-econometric modelling.

1.3 Overview

The remainder of the paper is structured as follows. Since we attribute a central role to forecast-period shifts in deterministic factors as the cause of forecast failure, section 2 first explains the concept of deterministic shifts, reviews their implications, and contrasts those with the impacts of non-deterministic shifts. Thereafter, the analysis assumes a world subject to such shifts. Section 3 derives the resulting implications for 'rational expectations', and suggests alternatives that are both feasible and more robust to breaks. Then section 4 discusses tests of theory-based propositions, before section 5 turns to model selection for forecasting. Next, section 6 introduces three related sections concerned with aspects of model selection for policy. First, section 6.1, considers the obverse of section 5, and shows that policy models should not be selected by forecast criteria. Secondly, section 6.2 considers policy analyses based on impulse responses, and thirdly, section 6.3 examines estimation of policy responses. Section 7 describes appropriate model-selection procedures, based on computer automation, and section 8 justifies the focus on congruent modelling. Finally, section 9 concludes.

2 Setting the scene

In a constant-parameter, stationary world, forecast failure should rarely occur: the in-sample and out-of-sample fits will be similar because the data properties are unchanged. As discussed in Miller (1978), stationarity ensures that, on average (i.e., excluding rare events), an incorrectly-specified model will forecast within its anticipated tolerances (providing these are correctly calculated). Although a mis-specified model could be beaten by methods based on correctly-specified equations, it will not suffer excessive forecast failure purely because it is mis-specified. Nevertheless, since a congruent, encompassing model will variance-dominate in-sample, it will continue to do so when forecasting under unchanged conditions. Thus, adding causal variables will improve forecasts on average; adding non-causal variables (i.e., variables that do not enter the DGP) will only do so when they proxy for omitted causal variables. In an important sense, the best model will win.

Empirical models are usually data-based (selected to match the available observations), which could induce some overfitting, but should not produce systematic forecast failure (see Clements and Hendry, 1999b). Conversely, when the data properties over the forecast horizon differ from those in-sample – a natural event in non-stationary processes – forecast failure will result. The latter’s regular occurrence is strong evidence for pandemic non-stationarity in economics, an unsurprising finding given the manifest legislative, social, technological and political changes witnessed over modern times (and indeed through most of history).

Once such non-stationarity is granted, many ‘conventional’ results that are provable in a constant-parameter, stationary setting change radically. In particular, since the future will not be like the present or the past, two important results can be established in theory, and demonstrated in practice: the potential forecast dominance of models using causal variables by those involving non-causal variables; and of in-sample well-specified models by badly mis-specified ones. Clements and Hendry (1999a) provide several examples: another is offered below. Together, such results remove the theoretical support for basing forecasting models – and hence agents’ expectations formation – on the in-sample conditional expectation given available information. We develop this analysis in section 3. Moreover, these two results potentially explain why over-differencing and intercept corrections – both of which introduce non-causal variables into forecasting devices – could add value to model-based forecasts: this aspect is explored in section 5, which emphasizes the potential dangers of selecting a policy model by such criteria as forecast accuracy. Finally, a failure to model the relevant non-stationarities can distort in-sample tests, and lead to incorrect inferences about the usefulness or otherwise of economic theories: that is the topic of section 4.

Not all forms of non-stationarity are equally pernicious. For example, unit roots generate stochastic trends in data series, which thereby have changing means and variances, but nevertheless seem relatively benign. This form of non-stationarity can be removed by differencing or cointegration transformations, and often, it may not matter greatly whether or not those transforms are imposed (see e.g., Sims, Stock and Watson, 1990, for estimation, and Clements and Hendry, 1998, for forecasting). Of course, omitting dynamics could induce ‘nonsense regressions’, but provided appropriate critical values are used, even that hypothesis is testable – and its rejection entails cointegration. As we show in section 2.2, shifts in parameters that do not produce any deterministic shifts also need not induce forecast failure, despite inducing non-stationarities.

Consider an h -step ahead forecast made at time T , denoted $\hat{\mathbf{y}}_{T+h|T}$, for a vector of n_y variables \mathbf{y}_{T+h} . The difference between the eventual outcomes and the forecast values is the vector of forecast errors $\mathbf{e}_{T+h|T} = \mathbf{y}_{T+h} - \hat{\mathbf{y}}_{T+h|T}$, and this can be decomposed into the various mistakes and unpredictable elements. Doing so delivers a forecast-error taxonomy, partitioned appropriately into deterministic, ob-

served stochastic, and innovation-error influences. For each component, there are effects from structural change, model mis-specification, data inaccuracy, and inappropriate estimation. Although the decomposition is not unique, it can be expressed in nearly-orthogonal effects corresponding to influences on forecast-error means and variances respectively. The former involves all the deterministic terms; the latter the remainder. We now briefly consider these major categories of error, commencing with mean effects, then turn to variance components.

2.1 Forecast failure and deterministic shifts

Systematic forecast-error biases derive from deterministic factors being mis-specified, mis-estimated, or non-constant. The simplest example is omitting a trend; or when a trend is included, under-estimating its slope; or when the slope is correct, experiencing a shift in the growth rate. A similar notion applies to equilibrium means, including shifts, mis-specification of, or mis-estimation in the means of (say) the savings rate, velocity of circulation, or the NAIRU. Any of these will lead to a systematic, and possibly increasing, divergence between outcomes and forecasts. However, there is an important distinction between the roles of intercepts, trends etc., in models, and any resulting deterministic shifts, as we will now explain.

To clarify the roles of deterministic, stochastic, and error factors, we consider a static regression where the parameters change prior to forecasting. The in-sample DGP, for $t = 1, \dots, T$, is:

$$y_t = \alpha + \beta x_t + \epsilon_t, \quad (1)$$

where x_t is an independent normally-distributed variable with mean μ and variance σ_x^2 . Also, ϵ_t is an independent, normally-distributed error with mean zero and constant variance σ_ϵ^2 . Finally, x_t has known future values to the forecaster, and $\{\epsilon_t\}$ is independent of x . A special case of interest below is $\beta = 0$, so α is just the mean of y .

In (1), the conditional mean of y_t is $E[y_t|x_t] = \alpha + \beta x_t$ and the conditional variance is $V[y_t] = \sigma_\epsilon^2$. The unconditional mean, $E[y_t]$, and variance, $V[y_t]$, (which allow for the variation in x_t) are $\alpha + \beta\mu = \gamma$ and $\sigma_\epsilon^2 + \beta^2\sigma_x^2$ respectively. Thus, there are two deterministic components in (1): the intercept α and the mean of the regressor μ , so the overall deterministic term is γ . Indeed, we can always rewrite (1) as:

$$\begin{aligned} y_t &= \alpha + \beta\mu + \beta(x_t - \mu) + \epsilon_t \\ &= \gamma + \beta(x_t - \mu) + \epsilon_t, \end{aligned} \quad (2)$$

Shifts in the composite deterministic term γ will transpire to be crucial.

Consider using the estimated DGP (1) as the forecasting model. For simplicity, we assume known parameter values. Then, with an exactly-measured forecast origin at time T , (1) produces the h -step ahead forecast sequence:

$$\hat{y}_{T+h|T} = \alpha + \beta x_{T+h}. \quad (3)$$

However, over the forecast period, $h = 1, \dots, H$, there is a shift in the parameters of the process unknown to the forecaster, so that in fact:

$$\begin{aligned} y_{T+h} &= \alpha^* + \beta^* x_{T+h} + \epsilon_{T+h} \\ &= \gamma^* + \beta^* (x_{T+h} - \mu) + \epsilon_{T+h}. \end{aligned}$$

The distributions of $\{x_{T+h}\}$ and $\{\epsilon_{T+h}\}$ could also change (e.g., μ, σ_x^2 or σ_ϵ^2 might change), but we neglect such effects here as not germane to the central issues. Indeed, when x_t is known, as is assumed here, changes μ in are irrelevant.

The resulting sequence of forecast errors $e_{T+h|T} = y_{T+h} - \hat{y}_{T+h|T}$ after the unanticipated shift is:

$$\begin{aligned} & \alpha^* + \beta^* x_{T+h} + \epsilon_{T+h} - (\alpha + \beta x_{T+h}) \\ = & (\alpha^* - \alpha) + (\beta^* - \beta) x_{T+h} + \epsilon_{T+h}. \end{aligned} \quad (4)$$

There are two kinds of terms in (4): those contributing to the mean, and deviations from that mean. The former is obtained by taking expectations, which leads to:

$$\begin{aligned} E_{T+h} [e_{T+h|T}] &= (\alpha^* - \alpha) + (\beta^* - \beta) \mu \\ &= \gamma^* - \gamma. \end{aligned} \quad (5)$$

The composite shift is zero if and only if $\gamma^* = \gamma$: importantly, that does not require $\alpha^* = \alpha$ and $(\beta^* - \beta) \mu = 0$. When (5) is non-zero, we call it a ‘deterministic shift’: the effect is pernicious when $E_{T+h} [e_{T+h|T}]$ increases by several σ_ϵ , but is then usually easy to detect.

An empirically-relevant case is when the variables labelled y_t and x_t are log-differences, so μ defines the mean growth rate of x_t , and γ that of y_t . Many econometric growth-rate equations have $\mu \simeq \sigma_\epsilon$ (often around 0.5%–1%), so the requirement that $\gamma^* - \gamma$ be as large as (say) $2\sigma_\epsilon$ is actually very strong: e.g., a *doubling* of the trend rate of growth. Consequently, even moderate trend shifts can be hard to detect till quite a few periods have elapsed.

To illustrate that an ‘incorrect’ model can outperform the in-sample DGP in forecasting, we return to the special case of (1) when $\beta^* = \beta = 0$. The expected forecast error sequence from using the in-sample DGP will be $(\alpha^* - \alpha)$. That remains true when the forecast origin moves through time to $T + 1$, $T + 2$ etc.: because the forecasting model remains unchanged, so do the average forecast errors. Consider, instead, using the naive predictor $\tilde{y}_{T+h|T+1} = y_{T+1}$. The resulting sequence of forecast errors will be similar to $e_{T+h|T}$ when the origin is T : unanticipated shifts after forecasting are bound to harm all methods. However, when forecasting from time $T + 1$ onwards, a different result ensues for $\tilde{e}_{T+h|T+1} = y_{T+h} - \tilde{y}_{T+h|T+1}$ because:

$$\begin{aligned} \tilde{e}_{T+h|T+1} &= y_{T+h} - y_{T+1} \\ &= \alpha^* + \epsilon_{T+h} - \alpha^* - \epsilon_{T+1} \\ &= \Delta_{h-1} \epsilon_{T+h}, \end{aligned} \quad (6)$$

where $\Delta_{h-1} \epsilon_{T+h} = \epsilon_{T+h} - \epsilon_{T+1}$. The last line of (6) has a mean of zero, despite the deterministic shift. Thus, on a bias criterion, $\tilde{y}_{T+h|T+1}$ outperforms the in-sample DGP (and could win on mean-square error), despite the fact that y_{t-1} is not a causal variable. Dynamics make the picture more complicated, but similar principles apply.

An interesting, and much studied, example of a deterministic shift concerns forecast failure in a model of narrow money (M1) in the UK after the Banking Act of 1984, which permitted interest payments on current accounts in exchange for all interest payments being after the deduction of ‘standard rate’ tax. The own rate of interest (R_o) changed from zero to near the value of the competitive rate (R_c : about 12 per cent per annum at the time) in about 6 quarters, inducing very large inflows to M1. Thus, a large shift occurred in the mean opportunity cost of holding money, namely a deterministic shift from R_c to $R_c - R_o$. Pre-existing models of M1 – which used the outside rate of interest R_c as the measure of opportunity cost – suffered marked forecast failure, which persisted for many years after the break. Models that correctly re-measured the opportunity cost by $R_c - R_o$ continued to forecast well, once the break was observed, and indeed had the same estimated parameter values after the break as before. However, methods analogous to $\tilde{y}_{T+h|T}$ also did not suffer forecast failure: see Clements and Hendry (1999c) for details and references.

In general, the key drivers of forecast failure are mis-specification of, uncertainty in, or changes to the conditional expectation (where that exists) given the history of the process. The mean forecast can differ from the correct conditional expectation due to biased estimation of the mean, or when there are unexpected shifts. Because forecast failure is usually judged relative to in-sample behaviour, the latter is the dominant cause. However, mis-estimation of coefficients of deterministic terms could be deleterious to forecast accuracy if estimation errors are large by chance.

2.2 Non-deterministic shifts

Having extracted the deterministic terms, all other factors fall under the heading of non-deterministic. The converse problem now occurs. Shifts in the coefficients of zero-mean variables have a surprisingly small impact on forecasts (as measured by the inability of parameter-constancy tests to detect the break). There are three consequences. First, such shifts seem an unlikely explanation for observed forecast failure. Secondly, changes in reaction parameters such as β are difficult to detect unless they induce a deterministic shift in the model, which cannot occur when x_t has mean zero ($\mu = 0$). This finding helps explain the absence of empirical evidence on the Lucas (1976) critique, as discussed in section 3. Finally, although they have relatively benign effects in the context of forecasting, undetected changes in reaction parameters could have disastrous effects on policy analyses – but we leave that story till section 6.2.

More formally, when x_t and y_t both have mean zero in (1), all terms in (4) have zero expectations, so no forecast bias results when β changes. There is an increase in the forecast error variance, from σ_ϵ^2 to $(\beta^* - \beta)^2 \sigma_x^2 + \sigma_\epsilon^2$, and the detectability (or otherwise) of the break depends on how much the variance increases. For $(\beta^* - \beta)^2 = 4\sigma_\epsilon^2$ (say), then the ratio is $1 + 4\sigma_x^2$ which can be difficult to detect against the background noise (see e.g., Hendry and Doornik, 1997, for simulation illustrations).

In a stationary dynamic process, an intercept like α also differs from the unconditional mean γ , and it is shifts in the latter which are again relevant to forecast failure. A strong, and corroborated, prediction is that shifts in both the intercepts and the regression parameters which leave the unconditional mean unchanged will not induce forecast failure, and tests will be relatively powerless to detect that anything has changed. For example, $\gamma^* = \gamma$ when $\alpha^* + \beta^*\mu = \alpha + \beta\mu$, even though every parameter has altered. Indeed, the situation where the unconditional mean is constant is precisely the same as when all the means are zero: Hendry and Doornik (1997) and section 6.2 below provide the details, and lead to the conclusions that shifts in unconditional means are a primary source of forecast failure, and other ‘problems’ are less relevant to forecast failure.

For example, omitting zero-mean stochastic components is unlikely to be a major source of forecast failure, but could precipitate failure if stochastic mis-specification resulted in deterministic shifts elsewhere in the economy affecting the model. Equally, the false inclusion of zero-mean stochastic variables is a secondary problem, whereas wrongly including regressors which experienced deterministic shifts could have a marked impact on forecast failure as the model mean shifts although the data mean does not.

Estimation uncertainty in the parameters of stochastic variables also seems to be a secondary problem, as such errors add variance terms of $O(1/T)$ for stationary components. Neither collinearity nor a lack of parsimony *per se* seem likely culprits, although interacting with breaks occurring elsewhere in the economy could induce problems.

Finally, better-fitting models have smaller error accumulation, but little can be done otherwise about forecast inaccuracy from that source.

2.3 Digression: modelling deterministic terms

Over long runs of historical time, all aspects of economic behaviour are probably stochastic. However, in shorter periods some variables may exhibit little variation (or deviation from trend) and so be well represented by a deterministic variable. Further, such variables might be subject to level shifts characterizing different epochs (see Anderson and Mizon, 1989). ‘Models’ of intercept shifts are easily envisaged, where they become drawings from a ‘meta’ distribution; or where large shocks persist but small do not; or they are functions of more basic causes – endogenous growth theory could be seen as one attempt to model the intercepts in growth-rate equations. Such re-representations do not alter the forecasting implications drawn above, merely re-interpret what we call deterministic shifts: the key issue is whether the average draw over the relevant horizon is close to that over the sample used, or differs therefrom. The latter induces forecast failure.

That concludes the ‘scene setting’ analysis, summarized as: deterministic shifts of the data relative to the model are the primary source of forecast failure. Monte Carlo evidence presented in several papers bears out the analytics: parameter non-constancy and forecast-failure tests reject for small changes in unconditional means, but not for substantial changes in dynamics, or in all parameters when that leaves equilibrium means unaltered (all measured as a proportion of σ_ϵ).

3 ‘Rational expectations’

When unanticipated deterministic shifts make an economy non-stationary, the formation of ‘rational expectations’ requires agents to know:

- all the relevant information;
- how every component enters the joint data density;
- the changes in that density at each point in time.

In terms of our scalar example, the model forecast error $e_{T+h|T}$ in (4) equals the ‘rational expectations’ error ϵ_{T+h} if and only if every other term is zero. Yet most shifts, and many of their consequences, cannot be anticipated: assuming knowledge of current and *future* deterministic shifts is untenable. Otherwise, the resulting forecasting device can be dominated by methods which use no causally-relevant variables. Thus, it ceases to be rational to try and form expectations using the current conditional expectation when that will neither hold in the relevant future, nor forecast more accurately than other devices. Agents will learn that they do better forming expectations from ‘robust forecasting rules’ – which adapt rapidly to deterministic shifts. These may provide an example of ‘economically-rational expectations’ as suggested by Feige and Pearce (1976), equating the marginal costs and benefits of improvements in the accuracy of expectations: Hendry (2000b) provides a more comprehensive discussion.

Robust forecasting rules need not alter with changes in policy. Of course, if agents fully understood a policy change and its implications, they would undoubtedly be able to forecast better: but that would require the benefits of doing so to exceed the costs. The problem for agents is compounded by the fact that many major policy changes occur in turbulent times, precisely when it is most difficult to form ‘rational expectations’, and when robust predictors may outperform. Thus, many agents may adopt the adaptive rules discussed above, consistent with the lack of empirical evidence in favour of the Lucas (1976) critique reported in Ericsson and Irons (1995). Consequently, if an econometric model used Δx_t as a replacement for the expected change $\Delta x_{t+1|t}^e$ when agents used robust rules, then the model’s parameters need not change even after forecast failure occurred. Alternatively, the unimportant consequences for forecasting of changes in reaction coefficients, rather than their absence, could account

for the lack of empirical evidence that the critique occurs, but anyway reduces its relevance. Hence, though it might be sensible to use ‘rational expectations’ for a congruent and encompassing model in a stationary world, in practice the evident non-stationarities make it inadvisable.

4 Model selection for theory testing

Although not normally perceived as a ‘selection’ issue, tests of economic theories based on whole-sample goodness of fit comparisons involve selection, and can be seriously misled by deterministic shifts. Three examples affected by unmodelled shifts are: lagged information from other variables appearing irrelevant, affecting tests of Euler equation theories; cointegration failing so long-run relationships receive no empirical support; and tests of forecast efficiency rejecting because of residual serial correlation induced *ex post* by an unpredictable deterministic shift. We address these in turn.

The first two are closely related, so our illustration concerns tests of the implications of the Hall (1978) Euler-equation consumption theory when credit rationing changes, as happened in the UK (see Muellbauer, 1994). The log of real consumers’ expenditure on non-durables and services (c) is not cointegrated with the log of real personal disposable income (y) over 1962(2)–1992(4): a unit-root test using 5 lags of each variable, a constant and seasonals delivers $t_{ur} = 0.97$ so does not reject (see Banerjee and Hendry, 1992, and Ericsson and MacKinnon, 1999, on the properties of this test). Nevertheless, the solved long-run relation is:

$$c = - \underset{(0.99)}{0.53} + \underset{(0.10)}{0.98} y + \textit{Seasonals}. \quad (7)$$

Lagged income terms are individually ($max\ t = 1.5$) and jointly ($F(5, 109) = 1.5$) insignificant in explaining $\Delta_4 c_t = c_t - c_{t-4}$. Such evidence appears to support the Hall life-cycle model, which entails that consumption changes are unpredictable, with permanent consumption proportional to fully-anticipated permanent income. As fig. 1a shows for annual changes, the data behaviour is at odds with the theory after 1985, since consumption first grows faster than income for several years, then falls faster – far from smoothing. Moreover, the large departure from equilibrium in (7) is manifest in panel b, resulting in a marked deterioration in the resulting (fixed-parameter) 1-step forecast errors from the model in Davidson, Hendry, Srba and Yeo (1978) after 1984(4) (the period to the right of the vertical line in fig. 1c). Finally, an autoregressive model for $\Delta\Delta_4 c_t = \Delta_4 c_t - \Delta_4 c_{t-1}$ produces 1-step forecast errors which are smaller than average after 1984(4): see panel (d). Such a result is consistent with a deterministic shift around the mid 1980s (see Hendry, 1994, and Muellbauer, 1994, for explanations based on financial deregulation inducing a major reduction in credit rationing), which neither precludes the *ex ante* predictability of consumption from a congruent model, nor consumption and income being cointegrated. The apparent insignificance of additional variables may be an artefact of mis-specifying a crucial shift, so the ‘selected’ model is not valid support for the theory. Conversely, non-causal proxies for the break may seem significant. Thus, models used to test theories should first be demonstrated to be congruent and encompassing.

We must stress that our example is not an argument against econometric modelling. While $\Delta\Delta_4 c_{t-1}$ may be a more robust forecasting device than the models extant at the time, it is possible in principle that the appropriate structural model – which built in changes in credit markets – would both have produced better forecasts and certainly better policy. For example, by 1985, building society data suggested that mortgages were available on much easier terms than had been the case historically, and ‘housing-equity withdrawal’ was already causing concern to policy makers. Rather, we are criticizing the practice of

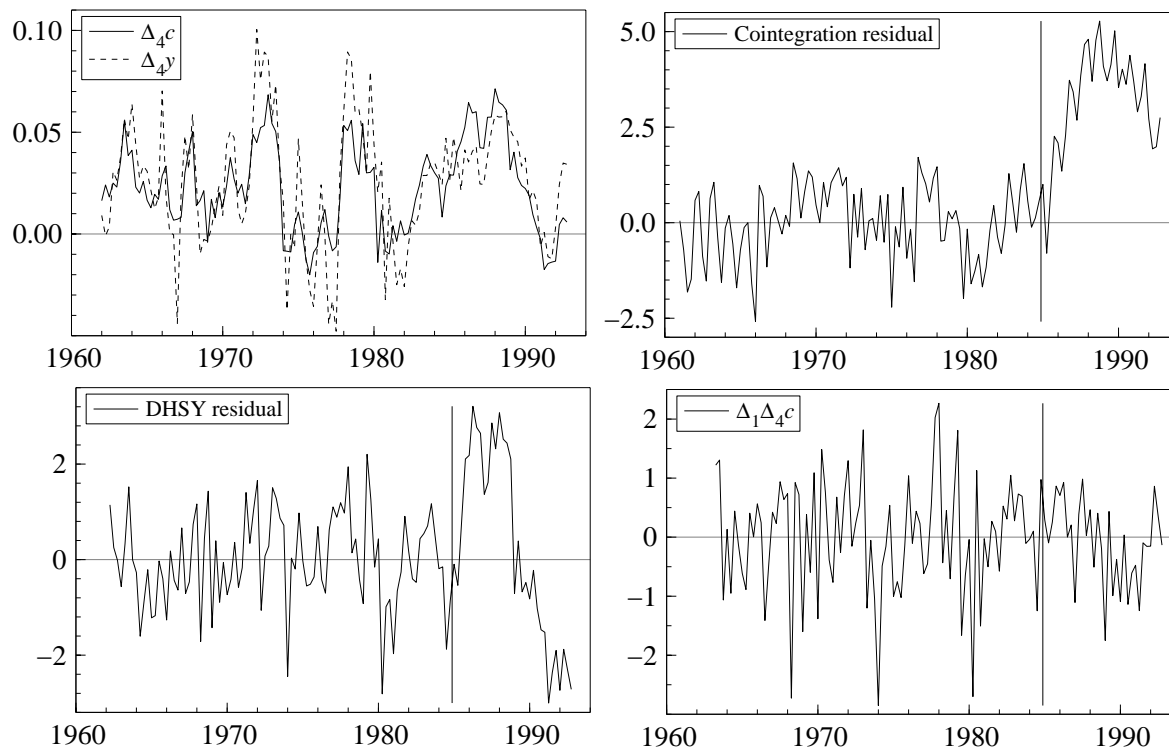


Figure 1 UK real consumers' expenditure and income with model residuals.

'testing theories' without first testing that the model used is a congruent and undominated representation, precisely because 'false' but robust predictors exist, and deterministic shifts appear to occur intermittently.

The same data illustrate the third mistake: rejecting forecast efficiency because of residual serial correlation induced *ex post* by an unpredictable deterministic shift. A model estimated prior to such a shift could efficiently exploit all available information; but if a shift was unanticipated *ex ante*, and unmodelled *ex post*, it would induce whole-sample residual serial correlation, apparently rejecting forecast efficiency. Of course, the results correctly reject 'no mis-specification'; but as no-one could have outperformed the in-sample DGP without prescience, the announced forecasts were not *ex ante* inefficient in any reasonable sense.

5 Model selection for forecasting

Forecast performance in a world of deterministic shifts is not a good guide to model choice, unless the sole objective is short-term forecasting. This is because models which omit causal factors and cointegrating relations, by imposing additional unit roots, may adapt more quickly in the face of unmodelled shifts, and so provide more accurate forecasts after breaks. We referred to this above as 'robustness' to breaks.

The admissible deductions on observing either the presence or absence of forecast failure are rather stark, particularly for general methodologies which believe that forecasts are the appropriate way to judge empirical models. In this setting of structural change, there may exist non-causal models (i.e., models none of whose 'explanatory' variables enter the DGP) that do not suffer forecast failure, and indeed may forecast absolutely more accurately on reasonable measures, than previously congruent, theory-based models. Conversely, *ex ante* forecast failure may merely reflect inappropriate measures

of the inputs, as we showed with the example of ‘opportunity-cost’ affecting UK M1: a model that suffers severe forecast failure may nonetheless have constant parameters on *ex post* re-estimation. Consequently, neither relative success nor failure in forecasting is a reliable basis for selecting between models – other than for forecasting purposes. Apparent failure on forecasting need have no implications for the goodness of a model, nor its theoretical underpinnings, as it may arise from incorrect data, that are later corrected.

Some forecast failures will be due to model mis-specification, such as omitting a variable whose mean alters; and some successes to having well-specified models that are robust to breaks. The problem is discriminating between such cases, since the event of success or failure *per se* is insufficient information. Because the future can differ in unanticipated ways from the past in non-stationary processes, previous success (failure) does not entail the same will be repeated later. That is why we have stressed the need for ‘robust’ or adaptable devices in the forecasting context. If it is desired to use a ‘structural’ or econometric model for forecasting, then there are many ways of increasing its robustness, as discussed in Clements and Hendry (1999a). The most usual are ‘intercept corrections’, which adjust the fit at the forecast origin to exactly match the data, and thereby induce the differences of the forecast errors that would otherwise have occurred. Such an outcome is close to that achieved by modelling the differenced data, but retains the important influences from disequilibria between levels. Alternatively, and closely related, one could update the equilibrium means and growth rates every period, placing considerable weight on the most recent data, retaining the in-sample values of all other reaction parameters.

In both cases, howsoever the adjustments are implemented, the policy implications of the underlying model are unaltered, although the forecasts may be greatly improved after deterministic shifts. The obvious conclusion, discussed further below, is that forecast performance is also not a good guide to policy-model choice. Without the correction, the forecasts would be poor; with the correction they are fine, but the policy recommendation is unaffected. Conversely, simple time-series predictors like $\Delta\Delta_4 c_{t-1}$ have no policy implications. We conclude that policy-analysis models should be selected on different criteria, which we now discuss.

6 Model selection for policy analysis

The next three sub-sections consider related, but distinct, selection issues when the purpose of modelling is policy: using forecast performance to select a policy model; investigating policy in closed models, where every variable is endogenous; and analyzing policy in open models, which condition on some policy variables. Although the three issues arise under the general heading of selecting a policy model, and all derive from the existence of, and pernicious consequences from, deterministic shifts, very different arguments apply to each, as we now sketch. ‘Selection’ is used in a general sense: only the first topic concerns an ‘empirical criterion’ determining the choice of model, whereas the other two issues derive from ‘prior’ decisions to select from within particular model classes.

First, because forecast failure derives primarily from unanticipated deterministic shifts, its occurrence does not sustain the rejection of a policy model: shifts in means may be pernicious, but need not impugn policy implications. For example, intercept corrections would have altered the forecast performance, but not the policy advice. Secondly, because badly mis-specified models can win forecasting competitions, forecast performance is not a sensible criterion for selecting policy models, as shown in section 6.1. Thirdly, policy conclusions depend on the values of reaction parameters, but we have noted the difficulty of detecting shifts in those parameters when there are no concomitant deterministic shifts, with adverse consequences for impulse-response analyses. Section 6.2 provides some detailed evidence.

Finally, policy changes in open models almost inevitably induce regime shifts with deterministic effects, and so can highlight previously hidden mis-specifications, or non-deterministic shifts; but also have a sustainable basis in the important concept of co-breaking, noted in section 1.2 above.

6.1 Selecting policy models by forecast performance

A statistical forecasting system is one having no economic-theory basis, in contrast to econometric models for which economic theory is the hallmark. Since the former system will rarely have implications for economic-policy analysis – and may not even entail links between target variables and policy instruments – being the ‘best’ available forecasting device is insufficient to ensure any value for policy analysis. Consequently, the main issue is the converse: does the existence of a dominating forecasting procedure invalidate the use of an econometric model for policy? Since forecast failure often results from factors unrelated to the policy change in question, an econometric model may continue to characterize the responses of the economy to a policy, despite its forecast inaccuracy.

Moreover, as stressed above, while such ‘tricks’ as intercept corrections may mitigate forecast failure, they do not alter the reliability of the policy implications of the resulting models. Thus, neither direction of evaluation is reliable: from forecast failure or success to poor or good policy advice. Policy models require evaluation on policy criteria.

Nevertheless, post-forecasting policy changes that entail regime shifts should induce breaks in models that do not embody the relevant policy links. Statistical forecasting devices will perform worse in such a setting: their forecasts are unaltered (since they do not embody the instruments), but the outcomes change. Conversely, econometric systems that do embody policy reactions need not experience any policy-regime shifts. Consequently, when both structural breaks and regime shifts occur, neither econometric nor time-series models alone are adequate: this suggests that they should be combined, and Hendry and Mizon (2000) provide an empirical illustration of doing so.

6.2 Impulse-response analyses

Impulse response analysis is a widely-used method for evaluating the response of one set of variables to ‘shocks’ in another set of variables (see e.g., Lütkepohl, 1991, Runkle, 1987, and Sims, 1980). The finding that shifts in the parameters of dynamic reactions are not readily detectable is potentially disastrous for impulse-response analyses of economic policy based on closed systems, usually vector autoregressions (VARs). Since changes in VAR intercepts and dynamic coefficient matrices may not be detected – even when tested for – but the full-sample estimates are a weighted average across different regimes, the resulting impulse responses need not represent the policy outcomes that will in fact occur. Indeed, this problem may be exacerbated by specifying VARs in first differences (as often occurs), since deterministic factors play a small role in such models.

It may be felt to be a cruel twist of fate that when a class of breaks is not pernicious for forecasting, it should be detrimental to policy – but these are just the opposite sides of the same coin. Moreover, this is only one of a sequence of drawbacks to using impulse responses on models to evaluate policy we have emphasized over recent years: see Banerjee, Hendry and Mizon (1996), Ericsson, Hendry and Mizon (1998a), and Hendry and Mizon (1998). Impulse response functions describe the dynamic properties of an estimated model, and not the dynamic characteristics of the variables. For example, when the DGP is a multivariate random walk, the impulse responses calculated from an estimated VAR in levels will rarely reveal the ‘persistence’ of shocks, since the estimated roots will not be exactly unity. Equally, estimated parameters may be inconsistent or inefficient, unless the model is congruent, encompasses rival models, and is invariant to extensions of the information used (see section 8). When a model has

the three properties just noted, it may embody structure (see Hendry, 1995b), but that does not imply that its residuals are structural: indeed, residuals cannot be invariant to extensions of information unless the model coincides with the DGP. In particular, increasing or reducing the number of variables directly affects the residuals, as does conditioning on putative exogenous variables. Worse still, specifying a variable to be weakly or strongly exogenous alters the impulse responses, irrespective of whether or not that variable actually is exogenous. While Granger non-causality (Granger, 1969) is sufficient for the equivalence of standard-error based impulse responses from systems and conditional models, it does not ensure efficient or valid inferences unless the conditioning variables are weakly exogenous. Moreover, the results are invariant to the ordering of variables only by ignoring the correlations between residuals in different equations. Avoiding this last problem by reporting orthogonalized impulse responses is not recommended either: it violates weak exogeneity for most orderings, induces a sequential conditioning of variables that depends on the chance ordering of the variables, and may lose invariance. The literature on ‘structural VARs’ (see, e.g., Bernanke, 1986, and Blanchard and Quah, 1989), which also analyzes impulse responses for a transformed system, faces a similar difficulty. The lack of understanding of the crucial role of weak exogeneity in impulse-response analyses is puzzling in view of the obvious feature that any given ‘shock’ to the error and to the intercept are indistinguishable, yet the actual reaction in the economy will be the same only if the means and variances are linked in the same way – which is the weak exogeneity condition in the present setting. Finally in closed systems which ‘model’ policy variables, impulse-response analysis assumes that the instrument process remains constant under the ‘shock’, when in fact this will often not be so. Thus, it may not be the response of agents that changes when there is change in policy, but via the policy feedback, the VAR coefficients themselves nevertheless shift, albeit in a way that is difficult to detect. There seems no alternative for viable policy analyses to carefully investigating the weak and super exogeneity status of appropriate policy conditioning variables.

Despite the fact that all of these serious problems are well known, impulse responses are still calculated. However, the problem which we are highlighting here – of undetectable breaks – is not well known, so we will demonstrate its deleterious impact using a Monte Carlo simulation. Consider the unrestricted $I(0)$ VAR:

$$\begin{aligned} y_{1,t} &= \phi_{11}y_{1,t-1} + \phi_{12}y_{2,t-1} + \epsilon_{1,t} \\ y_{2,t} &= \phi_{21}y_{1,t-1} + \phi_{22}y_{2,t-1} + \epsilon_{2,t} \end{aligned}$$

where both errors $\epsilon_{i,t}$ are independent, normally-distributed with means of zero and constant variances σ_{ii} , with $\mathbf{E}[\epsilon_{1,t}\epsilon_{2,s}] = 0 \forall t, s$. The $y_{i,t}$ are to be interpreted as $I(0)$ transformations of integrated variables, either by differencing or cointegration. We consider breaks in the $\Phi = (\phi_j)$ matrix, maintaining constant unconditional expectations of zero ($\mathbf{E}[y_{i,t}] = 0$). The full-sample size is $T = 120$, with a single break at $t = 0.5T = 60$, setting $\sigma_{ii} = 0.01$ (1% in a log-linear model). An unrestricted VAR with intercepts and one lag is estimated, and then tested for breaks. The critical values for the constancy tests are those for a *known* break point, which delivers the highest possible power for the test used. We consider a large parameter shift, from:

$$\Phi = \begin{pmatrix} 0.50 & -0.20 \\ -0.20 & -0.25 \end{pmatrix}, \quad (8)$$

to:

$$\Phi^* = \begin{pmatrix} 0.50 & 0.20 \\ 0.20 & 0.25 \end{pmatrix}, \quad (9)$$

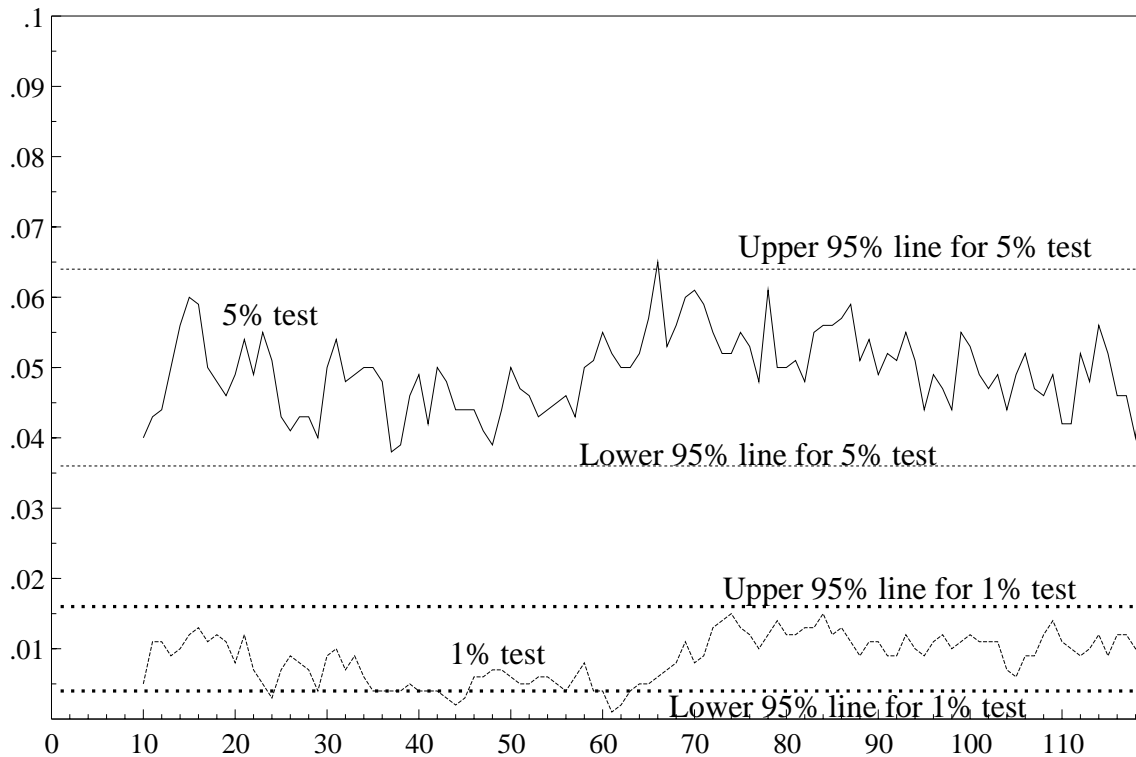


Figure 2 Constancy-test rejection frequencies for the $l(0)$ null.

so the sign is altered on all but one response, which is left constant simply to highlight the changes in the other impulses below.

We computed 1000 replications at both $p = 0.05$ and $p = 0.01$ (the estimates have standard errors of about 0.007 and 0.003 respectively), both when the null is true (no break) and when the break from Φ to Φ^* occurs. The resulting constancy-test rejection frequencies are reported graphically for both p values to illustrate the outcomes visually: the vertical axes show the rejection frequencies plotted against the sample sizes on the horizontal.

6.2.1 Test rejection frequencies under the null

As fig. 2 reveals, the null rejection frequencies in the $l(0)$ baseline data are reassuring: with 1000 replications, the approximate 95% confidence intervals are (0.036, 0.064) and (0.004, 0.016) for 5% and 1% nominal, and these are shown on the graphs as dotted and dashed lines respectively. The actual test null rejection frequencies are, therefore, close to their nominal levels. This gives us confidence that the estimated power to detect the break is reliable.

6.2.2 Shift in the dynamics

The constancy-test graph in fig. 3 shows the rejection frequencies when a break occurs. The highest power is less than 25%, even though the change constitutes a major structural break for the model economy: the detectability of a shift in dynamics is low when the DGP is an $l(0)$ VAR. This may be an explanation for the lack of evidence supporting the Lucas (1976) critique: shifts in zero-mean reaction parameters are relatively undetectable, rather than absent.




Figure 3 Constancy-test rejection frequencies for the $l(0)$ structural break.

6.2.3 Misleading impulse responses

Finally, we record the impulse responses from the averages of pre- and post- break models, and the model fitted across the regime shifts in fig. 4. The contrast is marked: despite the near undetectability of the break, the signs of most of the impulses have altered, and those obtained from the fitted model sometimes reflect one regime, and sometimes the other. Overall, mis-leading policy advice would follow, since even testing for the break would rarely detect it.

6.3 Policy analysis in open-models

Many of the problems in analyzing the responses of targets to changes in instruments noted above are absent when the modelling is validly conditional on the instruments, leading to an open model. Since it is often difficult to model the time-series behaviour of the policy instruments, particularly in high-dimensional systems, conditioning on them is much easier and is certainly preferable to omitting them from the analysis. For economic policy analysis, another advantage of modelling n_y target variables y_t conditionally on n_z instrument variables z_t , is that $\partial y_{t+h} / \partial z'_t$, which are important ingredients in the required policy responses, are directly estimable analytically, or at worst via simulation.

The fact that the $\{z_t\}$ process is under the control of a policy agency does not ensure that z_t are exogenous variables: indeed, policy is likely to depend on precisely the disequilibria in the rest of the economy that are key to its internal dynamics. Although the weak exogeneity of z_t for the parameters of the endogenous variables' equations is required for there to be no loss of information in making inferences on the parameters of interest, in practice it is likely that reliable estimates of policy responses will be obtained even when z_t is not weakly exogenous. This is because cointegration relations are usually established in the system, before conditioning is introduced.

A more important requirement is that whenever policy involves a regime shift, the instruments must be super exogenous for the parameters of interest. Co-breaking (described in section 1.2 above) between




Figure 4 Impulse response comparisons in an $I(0)$ VAR.

the targets and instruments then ensures that the policy is effective, and that the response of \mathbf{y}_t can be reliably estimated (efficiently, when \mathbf{z}_t is weakly exogenous for the response parameters). Since realistic policies involve deterministic shifts, any failure of co-breaking will be readily detected.

6.3.1 Stationary process

This section draws on some results in Ericsson *et al.* (1998a), and is presented for completeness as a preliminary to considering the more realistic integrated setting in the next sub-section. Modelling the conditional distribution for \mathbf{y}_t given \mathbf{z}_t (and any relevant lags) will yield efficient inference on the parameters of interest when \mathbf{z}_t is weakly exogenous for those parameters. In addition, the conditional model will provide reliable estimates of the response in \mathbf{y}_t to policy changes in \mathbf{z}_t when its parameters are invariant to the policy change. When these conditions are satisfied, the conditional model provides viable impulse responses and dynamic multipliers for assessing the effects of policy. However, Ericsson *et al.* (1998a) showed that, in general, the weak exogeneity status of conditioning variables is not invariant to transformations such as orthogonalizations, or identified ‘structural VARs’.

Irrespective of the exogeneity status of \mathbf{z}_t , modelling the conditional distribution alone will result in different impulse-response matrices ($\partial \mathbf{y}_{t+h} / \partial \boldsymbol{\varepsilon}'_t$) and dynamic multipliers ($\partial \mathbf{y}_{t+h} / \partial \mathbf{z}'_t$) because the latter takes into account the effects from contemporaneous and lagged \mathbf{z}_t . Thus, the response of \mathbf{y}_t to an impulse in the innovation $\boldsymbol{\varepsilon}_t$ of the conditional model is not the relevant response for assessing the effects of policy changes in the $\{\mathbf{z}_t\}$ process.

6.3.2 Integrated process

In an integrated process, the class of open models will be equilibrium-correction systems conditional on the current growth-rate of the policy instruments, $\Delta \mathbf{z}_t$, assuming that the \mathbf{z}_t are $I(1)$, and are included

in some of the r cointegration relations $\beta' \mathbf{x}_{t-1}$. For there to be no loss of information, this analysis requires that \mathbf{z}_t be weakly exogenous for both the long-run parameters β , and any short-run dynamic response parameters (for both lagged \mathbf{y}_{t-s} and lagged \mathbf{z}_{t-k}), all of which parameters should be invariant to the policy change. Under these conditions, it is possible to estimate the responses in the growth rates $\Delta \mathbf{y}_t$ and the disequilibria $\beta' \mathbf{x}_t$ to particular choices of the instruments \mathbf{z}_t even when the latter are $I(1)$. To derive the impact on (say) $\Delta \mathbf{y}_{t+h}$ from a change in \mathbf{z}_t requires a specification of the future path of \mathbf{z}_{t+i} in response to \mathbf{z}_t over $t = 1$ to $t + h$; implicitly, the model must be closed. This provides a link to the ‘policy rules literature’ (see e.g., Taylor, 1993, 2000), where alternative mappings of the policy instruments onto past values of disequilibria are evaluated. Nevertheless, the outcomes obtained can differ substantially from impulse-response analysis based on a (cointegrated) VAR when the policy rule does not coincide with the historical description of policy responses, and in importantly when the policy rule itself is changed perhaps as a result of observing previous responses to policy.

Partitioning the disequilibria $\beta' \mathbf{x}_t = \beta'_y \mathbf{y}_t + \beta'_z \mathbf{z}_t$ reveals that $\beta'_y \mathbf{y}_t$ are feasible target variables in this context despite \mathbf{y}_t and $\beta'_y \mathbf{y}_t$ being $I(1)$. However, an implication of this analysis is that very special conditions are required for policy that changes a single instrument $\tilde{z}_{i,t}$ (e.g., the minimum lending rate) to successfully target a single target variable $y_{j,t}$ (e.g., inflation) when these variables are $I(1)$. Conversely, Johansen and Juselius (2000) demonstrate that if a policy that targets an $I(1)$ variable is successful, then the target will be rendered $I(0)$.

Also, there are only $n_y + n_z - r$ unconstrained stochastic trends when r is equal to the number of cointegrating vectors, so the growth rates π_y of \mathbf{y}_t and π_z of \mathbf{z}_t are linked by $\beta'_y \pi_y + \beta'_z \pi_z = \mathbf{0}$. However, when there is co-breaking between $\Delta \mathbf{y}_t$ and $\Delta \mathbf{z}_t$, then a change in π_z will result in a corresponding change in the unconditional mean π_y of $\Delta \mathbf{y}_t$. Hence, just as in Hendry and Mizon (1998, 2000), linkages between deterministic terms are critical for policy to be effective when it is implemented via shifts in deterministic terms in the instrument process. Moreover, co-breaking here requires that $\Delta \mathbf{y}_t$ responds to contemporaneous and/or lagged changes in \mathbf{z}_t .

An important aspect of policy changes which comprise deterministic shifts is their ability to reveal previously undetected changes which might contaminate model specification. The dynamic response in a model will trace out a sequence of shifts over time, which will differ systematically from the corresponding responses in the economy when earlier changes lurk undetected. While the outcome will not be as anticipated, the mis-specification does not persist undetected.

7 Empirical model selection

First developing congruent general models, then selecting appropriate simplifications thereof that retain only the relevant information has not proved easy – even to experienced practitioners. The former remains the domain where considerable detailed institutional, historical and empirical knowledge interacts with value-added insights and clever theories of investigators: a good initial general model is essential. However, the latter is primarily determined by econometric modelling skills, and the developments in Hoover and Perez (1999) suggest automating those aspects of the task that require the implementation of selection rules, namely the simplification process. Just as early chess-playing programs were easily defeated, but later ones can systematically beat Grandmasters, so we anticipate computer-automated model selection software will develop well beyond the capabilities of the most expert modellers. We now explain why a general-to-specific modelling strategy – as implemented in *PcGets* – is able to perform so well despite the problem of ‘data mining’, discuss the costs of search, distinguish them from the (unavoidable) costs of inference, and suggest that the practical modelling problem is to retain relevant

variables, not eliminate spurious ones.

Statistical inference is always uncertain because of type I and type II errors (rejecting the null when it is true; and failing to reject the null when it is false respectively). Even if the DGP were derived *a priori* from economic theory, an investigator could not *know* that such a specification was ‘true’, and inferential mistakes will occur when testing hypotheses about it. This is a ‘pre-test’ problem: beginning with the truth and testing it will sometimes lead to false conclusions. ‘Pre-testing’ is known to bias estimated coefficients, and may distort inference (see *inter alia*, Judge and Bock, 1978). Of course, the DGP specification is never known in practice, and since ‘theory dependence’ in a model has as many drawbacks as ‘sample dependence’, data-based model-search procedures are used in practice, thus adding search costs to the costs of inference. A number of arguments point towards the advantages of ‘general-to-specific’ searches.

Statistical analyses of repeated testing provide a pessimistic background: every test has a non-zero null rejection frequency (‘size’), so type I errors accumulate. Size could be lowered by increasing the significance levels of selection tests, but only at the cost of reducing power to detect the influences that really matter. The simulation experiments in Lovell (1983) suggested that search had high costs, leading to an adverse view of ‘data mining’. However, he evaluated outcomes against the truth, compounding costs of inference with costs of search. Rather, the key issue for any model-selection procedure is: how costly is it to search across many alternatives relative to commencing from the DGP? As we now discuss, it is feasible to lower size and raise power simultaneously by improving the search algorithm.

First, White (1990) showed that with sufficiently-rigorous testing and a large enough data sample, the selected model will converge to the DGP, so selection error is a ‘small-sample’ problem, albeit a difficult and prevalent one. Secondly, Mayo (1981) noted that diagnostic testing was effectively independent of the sufficient statistics from which parameter estimates are derived, so would not distort the latter. Thirdly, since the DGP is obviously congruent with itself, congruent models are the appropriate class within which to search. This argues for commencing the search from a congruent model. Fourthly, encompassing – explaining the evidence relevant for all alternative models under consideration – resolves ‘data mining’ (see Hendry, 1995a) and delivers a dominant outcome. This suggests commencing from a general model that embeds all relevant contenders. Fifthly, any model-selection process must avoid getting stuck in search paths that inadvertently delete relevant variables, thereby retaining many other variables as proxies. The resulting approach of sequentially simplifying a congruent general unrestricted model (GUM) to obtain the maximal acceptable reduction, is called general-to-specific (*Gets*).

To evaluate the performance of *Gets* modelling procedures, Hoover and Perez (1999) reconsidered the Lovell (1983) experiments, searching for a single conditional equation (with 0 to 5 regressors) from a large macroeconomic database (containing up to 40 variables, including lags). By following several reduction search paths – each terminated by either no further feasible reductions or significant diagnostic test outcomes – they showed how much better the structured *Gets* approach was than any method Lovell considered, suggesting that modelling *per se* need not be bad. Indeed, the overall ‘size’ (false null rejection frequency) of their selection procedure was close to that expected without repeated testing, yet the power was reasonable.

Building on their findings, Hendry and Krolzig (1999) and Krolzig and Hendry (2000) developed the Ox (see Doornik, 1999) program *PcGets* which first tests the congruency of a GUM, then conducts pre-selection tests for ‘highly irrelevant’ variables at a loose significance level (25% or 50%, say), and simplifies the model accordingly. It then explores many selection paths to eliminate statistically-insignificant variables on F- and t-tests, applying diagnostic tests to check the validity of all reductions, thereby ensuring a congruent final model. All the terminal selections resulting from search paths are stored, and encompassing procedures and information criteria select between the contenders. Finally,

sub-sample significance is used to assess the reliability of the resulting model choice. In Monte Carlo experiments, *PcGets* recovers the DGP with power close to what one would expect if the DGP were known, and empirical size often below the nominal, suggesting that search costs are in fact low. In the ‘classic’ experiment in which the dependent variable is regressed on 40 irrelevant regressors, *PcGets* correctly finds the null model about 97% of the time for the Lovell database.

Some simple analytics proposed in Hendry (2000a) suggest why *PcGets* performs well, even though the following analysis ignores pre-selection, search paths and diagnostic testing (all of which improve the algorithm). An F-test against the GUM using critical value c_γ would have size $P(F \geq c_\gamma) = \gamma$ under the null if it were the only test implemented. For k regressors, the probability of retaining no variables from t-tests at size α is:

$$P(|t_i| < c_\alpha \forall i = 1, \dots, k) = (1 - \alpha)^k, \quad (10)$$

where the average number of variables retained then is:

$$n = k\alpha. \quad (11)$$

Combined with the F-test of the GUM, the probability π of correctly selecting the null model is no smaller than:

$$\pi = (1 - \gamma) + \gamma(1 - \alpha)^k. \quad (12)$$

For $\gamma = 0.05$ and $\alpha = 0.01$, when $k = 40$, then $\pi = 0.98$ and $n = 0.4$. Although falsely rejecting the null on the F-test signals that spurious significance lurks, so (11) will understate the number of regressors then retained, nevertheless, eliminating adventitiously-significant (spurious) variables is not the real problem in empirical modelling.

Indeed, the focus in earlier research on ‘over-fitting’ – reflecting inferior algorithms – has mis-directed the profession’s attention. The really difficult problem is retaining the variables that matter. Consider an equation with six relevant regressors, all with (absolute) t-values of 2 on average (i.e., $E[|t_i|] = 2$). The probability in any given sample that each observed $\widehat{t}_i \geq c_\alpha = 2$ (say) is approximately 0.5, so even if one began with the DGP, the probability of retaining all six is:

$$\begin{aligned} P(\widehat{t}_i \geq c_\alpha \forall i = 1, \dots, k \mid |t_i| = 2) &= 0.5^6 \\ &\simeq 0.016. \end{aligned}$$

Using 1% significance lowers this to essentially zero. Surprisingly, even if every $E[|t_i|] = 3$, the chances of keeping the DGP specification are poor:

$$P(\widehat{t}_i \geq c_\alpha \forall i \mid |t_i| = 3) = 0.84^6 \simeq 0.35.$$

Thus, the costs of inference are high in such full-sample testing, and will lead to under-estimating model size. An alternative, block-testing, approach discussed in Hendry (2000a) seems able to improve the power substantially.

Nevertheless, many empirical equations have many regressors. This is probably due to the high average t-values found in economics:

$$P(\widehat{t}_i \geq c_\alpha \forall i \mid |t_i| = 5) \simeq 0.989^6 \simeq 0.935,$$

(so almost all will always be retained), and not to selection biases as shown above. Even selecting by t-testing from 40 candidate regressors at 5% would only deliver 2 significant variables on average.

We conclude that models with many significant variables correctly represent some of the complexity of aggregate economic behaviour and not ‘over-fitting’.

The evidence to date in Hoover and Perez (1999) and Hendry and Krolzig (1999) for conditional dynamic models, Krolzig (2000) for VARs, and Hoover and Perez (2000) for cross-section data sets is of equally impressive performance by *Gets* in model selection. Certainly, the Monte Carlo evidence concerns selecting a model which is a special case of the DGP, whereas empirically, models are unlikely even to be special cases of the local DGP (LDGP). However, that is not an argument against *Gets*, but confirms the need for commencing with general specifications than have some chance of embedding the LDGP. Only then will reliable approximations to the actual economic process be obtained.

8 Congruent modelling

As a usable knowledge base, theory-related, congruent and encompassing econometric models remain undominated by matching the data in all measurable respects (see, e.g., Hendry, 1995a). For empirical understanding, such models seem likely to remain an integral component of any progressive research strategy. Nevertheless, even the ‘best available model’ can be caught out when forecasting by an unanticipated outbreak of (say) a major war or other crisis for which no effect was included in the forecast. However, if empirical models which are congruent within sample remain subject to a non-negligible probability of failing out of sample, then a critic might doubt their worth. Our defence of the program of attempting to discover such models rests on the fact that empirical research is part of a progressive strategy, in which knowledge gradually accumulates. This includes knowledge about general causes of structural changes, such that later models incorporate measures accounting for previous events, and hence are more robust (e.g., to wars, changes in credit rationing, financial innovations, etc.). For example, the dummy variables for purchase-tax changes in Davidson *et al.* (1978) that at the time ‘mopped up’ forecast failure, later successfully predicted the effects of introducing VAT, as well as the consequences of its doubling in 1979; and the First World-War shift in money demand in Ericsson, Hendry and Prestwich (1998b) matched that needed for the Second World War.

Since we now have an operational selection methodology with excellent properties, *Gets* seems a natural way to select models for empirical characterization, theory testing and policy analyses. When the GUM is a congruent representation, embedding the available theory knowledge of the target-instrument linkages, and parsimoniously encompassing previous empirical findings, the selection strategy described in section 7 offers scope for selecting policy models. Four features favour such a view. First, for a given null rejection frequency, variables that matter in the DGP are selected with the same probabilities as if the DGP were known. In the absence of omniscience, it is difficult to imagine doing much better systematically. Secondly, although estimates are biased on average, conditional on retaining a variable, its coefficient provides an unbiased estimate of the policy reaction parameter. This is essential for economic policy – if a variable is included, *PcGets* delivers the right response; otherwise, when it is excluded, one is simply unaware that such an effect exists.¹ Thirdly, the probability of retaining adventitiously significant variables is around the anticipated level for the variables that remain after pre-selection simplification. If that is (say) even as many as 30 regressors, of which 5 actually matter, then at 1% significance, 0.25 extra variables will be retained on average: i.e., one additional ‘spuriously-significant’ variable per four equations. This seems unlikely to distort policy in important ways. Finally,

¹This is one of three reasons why we have not explored ‘shrinkage’ estimators, which have been proposed as a solution to the ‘pre-test’ problem, namely, they deliver biased estimators (see, e.g., Judge and Bock, 1978). The second, and main, reason is that such a strategy has no theoretical underpinnings in processes subject to intermittent parameter shifts. The final reason concerns the need for progressivity, explaining more by less, which such an approach hardly facilitates.

the sub-sample – or more generally, recursive – selection procedures help to reveal which variables have non-central t-statistics, and which central (and hence should be eliminated). Overall, the role of *Gets* in selecting policy models looks promising.

Because changes to the coefficients of zero-mean variables are difficult to detect in dynamic models, for policy models they remain hazardous: the estimated parameters would appear to be constant, but in fact be mixtures across regimes, leading to inappropriate advice. In a progressive research context (i.e., from the perspective of learning), this is unproblematic since most policy changes involve deterministic shifts (as opposed to mean-preserving spreads), hence earlier incorrect inferences will be detected rapidly – but is cold comfort to the policy maker, or the economic agents subjected to the wrong policies.

9 Conclusion

The implications for econometric modelling that result from the observance of forecast failure differ considerably from those obtained when the model is assumed to coincide with a constant mechanism. Causal information can no longer be shown to uniformly dominate non-causal. Intercept corrections have no theoretical justification in stationary worlds with correctly-specified empirical models, but in a world subject to structural breaks of unknown form, size, and timing, they serve to ‘robustify’ forecasts against deterministic shifts – as the practical efficacy of intercept corrections confirms. Forecasting success is no better an index for model selection than forecast failure is for model rejection. Thus, emphasizing ‘out-of-sample’ forecast performance (perhaps because of fears over ‘data-mining’) is unsustainable (see, e.g., Newbold, 1993, p.658), as is the belief that a greater reliance on economic theory will help forecasting (see, e.g., Diebold, 1998), because that does not tackle the root problem.

A taxonomy of potential sources of forecast errors clarifies the roles of model mis-specification, sampling variability, error accumulation, forecast origin mis-measurement, intercept shifts, and slope-parameter changes. Forecast failure seems primarily attributable to deterministic shifts in the model relative to the data. Other shifts are far more difficult to detect. Such findings are potentially disastrous for ‘impulse-response’ analyses of economic policy. Since the changes in VAR intercepts and dynamic coefficient matrices may not be detected even when tested for, but the recorded estimates are a weighted average across the different regimes, the resulting impulse responses do not represent the policy outcomes that will in fact occur.

If the economy were reducible by transformations to a stationary stochastic process, where the resulting unconditional moments were constant over time, then well-tested, causally-relevant, congruent models which embodied valid theory restrictions would both fit best, and by encompassing, also dominate in forecasting on average. The prevalence historically of unanticipated deterministic shifts suggests that such transformations do not exist. Even the best policy model may fail at forecasting in such an environment. As we have shown, this need not impugn its policy relevance – other criteria than forecasting are needed for that judgement. Nevertheless, the case for continuing to use econometric systems probably depends on their competing reasonably successfully in the forecasting arena. Cointegration, co-breaking, and model-selection procedures as good as *PcGets*, with rigorous testing will help in understanding economic behaviour and evaluating policy options, but none of these ensures immunity to forecast failure from new breaks. An approach which incorporates causal information in a congruent econometric system for policy, but operates with robustified forecasts, merits consideration. We have not yet established formally that *Gets* should be used for selecting policy models from a theory-based GUM – but such a proof should be possible, given the relative accuracy with which the DGP is located. Achieving that aim represents the next step of our research program, and we anticipate establishing that

a data-based *Gets* approach will perform well in selecting models for policy.

References

- Anderson, G. J., and Mizon, G. E. (1989). What can statistics contribute to the analysis of economic structural change?. In Hackl, P. (ed.), *Statistical Analysis and Forecasting of Economic Structural Change*, Ch. 1. Berlin: Springer-Verlag.
- Banerjee, A., and Hendry, D. F. (1992). Testing integration and cointegration: An overview. *Oxford Bulletin of Economics and Statistics*, **54**, 225–255.
- Banerjee, A., Hendry, D. F., and Mizon, G. E. (1996). The econometric analysis of economic policy. *Oxford Bulletin of Economics and Statistics*, **58**, 573–600.
- Bernanke, B. S. (1986). Alternative explorations of the money-income correlation. In Brunner, K., and Meltzer, A. H. (eds.), *Real Business Cycles, Real Exchange Rates, and Actual Policies*, Vol. 25 of *Carnegie-Rochester Conferences on Public Policy*, pp. 49–99. Amsterdam: North-Holland Publishing Company.
- Blanchard, O., and Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review*, **79**, 655–673.
- Clements, M. P., and Hendry, D. F. (1998). *Forecasting Economic Time Series*. Cambridge: Cambridge University Press.
- Clements, M. P., and Hendry, D. F. (1999a). *Forecasting Non-stationary Economic Time Series*. Cambridge, Mass.: MIT Press.
- Clements, M. P., and Hendry, D. F. (1999b). Modelling methodology and forecast failure. Unpublished typescript, Economics Department, University of Oxford.
- Clements, M. P., and Hendry, D. F. (1999c). On winning forecasting competitions in economics. *Spanish Economic Review*, **1**, 123–160.
- Davidson, J. E. H., Hendry, D. F., Srba, F., and Yeo, J. S. (1978). Econometric modelling of the aggregate time-series relationship between consumers' expenditure and income in the United Kingdom. *Economic Journal*, **88**, 661–692. Reprinted in Hendry, D. F., *Econometrics: Alchemy or Science?* Oxford: Blackwell Publishers, 1993, and Oxford University Press, 2000.
- Diebold, F. X. (1998). The past, present and future of macroeconomic forecasting. *The Journal of Economic Perspectives*, **12**, 175–192.
- Doornik, J. A. (1999). *Object-Oriented Matrix Programming using Ox* 3rd edn. London: Timberlake Consultants Press.
- Ericsson, N. R. (1992). Cointegration, exogeneity and policy analysis: An overview. *Journal of Policy Modeling*, **14**, 251–280.
- Ericsson, N. R., Hendry, D. F., and Mizon, G. E. (1998a). Exogeneity, cointegration and economic policy analysis. *Journal of Business and Economic Statistics*, **16**, 370–387.
- Ericsson, N. R., Hendry, D. F., and Prestwich, K. M. (1998b). The demand for broad money in the United Kingdom, 1878–1993. *Scandinavian Journal of Economics*, **100**, 289–324.
- Ericsson, N. R., and Irons, J. S. (1995). The Lucas critique in practice: Theory without measurement. In Hoover, K. D. (ed.), *Macroeconometrics: Developments, Tensions and Prospects*. Dordrecht: Kluwer Academic Press.
- Ericsson, N. R., and MacKinnon, J. G. (1999). Distributions of error correction tests for cointegration.

- International finance discussion paper no. 655, Federal Reserve Board of Governors, Washington, D.C. www.bog.frb.fed.us/pubs/ifdp/1999/655/default.htm.
- Feige, E. L., and Pearce, D. K. (1976). Economically rational expectations. *Journal of Political Economy*, **84**, 499–522.
- Fildes, R. A., and Makridakis, S. (1995). The impact of empirical accuracy studies on time series analysis and forecasting. *International Statistical Review*, **63**, 289–308.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, **37**, 424–438.
- Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Evidence. *Journal of Political Economy*, **86**, 971–987.
- Hendry, D. F. (1994). HUS revisited. *Oxford Review of Economic Policy*, **10**, 86–106.
- Hendry, D. F. (1995a). *Dynamic Econometrics*. Oxford: Oxford University Press.
- Hendry, D. F. (1995b). Econometrics and business cycle empirics. *Economic Journal*, **105**, 1622–1636.
- Hendry, D. F. (2000a). *Econometrics: Alchemy or Science?* Oxford: Oxford University Press. New Edition.
- Hendry, D. F. (2000b). Forecast failure, expectations formation, and the Lucas critique. Mimeo, Nuffield College, Oxford.
- Hendry, D. F., and Doornik, J. A. (1997). The implications for econometric modelling of forecast failure. *Scottish Journal of Political Economy*, **44**, 437–461. Special Issue.
- Hendry, D. F., and Krolzig, H.-M. (1999). Improving on ‘Data mining reconsidered’ by K.D. Hoover and S.J. Perez. *Econometrics Journal*, **2**, 202–219.
- Hendry, D. F., and Mizon, G. E. (1998). Exogeneity, causality, and co-breaking in economic policy analysis of a small econometric model of money in the UK. *Empirical Economics*, **23**, 267–294.
- Hendry, D. F., and Mizon, G. E. (2000). On selecting policy analysis models by forecast accuracy. In Atkinson, A. B., Glennerster, H., and Stern, N. (eds.), *Putting Economics to Work: Volume in Honour of Michio Morishima*, pp. 71–113. London School of Economics: STICERD.
- Hendry, D. F., and Richard, J.-F. (1982). On the formulation of empirical models in dynamic econometrics. *Journal of Econometrics*, **20**, 3–33. Reprinted in Granger, C. W. J. (ed.) (1990), *Modelling Economic Series*. Oxford: Clarendon Press and in Hendry D. F. (1993, 2000), *op. cit.*
- Hoover, K. D., and Perez, S. J. (1999). Data mining reconsidered: Encompassing and the general-to-specific approach to specification search. *Econometrics Journal*, **2**, 167–191.
- Hoover, K. D., and Perez, S. J. (2000). Truth and robustness in cross-country growth regressions. unpublished paper, Economics Department, University of California, Davis.
- Johansen, S., and Juselius, K. (2000). How to control a target variable in the VAR model. Mimeo, European University of Institute, Florence.
- Judge, G. G., and Bock, M. E. (1978). *The Statistical Implications of Pre-Test and Stein-Rule Estimators in Econometrics*. Amsterdam: North Holland Publishing Company.
- Krolzig, H.-M. (2000). General-to-specific reductions of vector autoregressive processes. Unpublished paper, Department of Economics, University of Oxford.
- Krolzig, H.-M., and Hendry, D. F. (2000). Computer automation of general-to-specific model selection procedures. *Journal of Economic Dynamics and Control*. forthcoming.
- Lovell, M. C. (1983). Data mining. *Review of Economics and Statistics*, **65**, 1–12.

- Lucas, R. E. (1976). Econometric policy evaluation: A critique. In Brunner, K., and Meltzer, A. (eds.), *The Phillips Curve and Labor Markets*, Vol. 1 of *Carnegie-Rochester Conferences on Public Policy*, pp. 19–46. Amsterdam: North-Holland Publishing Company.
- Lütkepohl, H. (1991). *Introduction to Multiple Time Series Analysis*. New York: Springer-Verlag.
- Makridakis, S., and Hibon, M. (2000). The M3-competition: Results, conclusions and implications. Discussion paper, INSEAD, Paris.
- Mayo, D. (1981). Testing statistical testing. In Pitt, J. C. (ed.), *Philosophy in Economics*, pp. 175–230: D. Reidel Publishing Co. Reprinted as pp. 45–73 in Caldwell B. J. (1993), *The Philosophy and Methodology of Economics*, Vol. 2, Aldershot: Edward Elgar.
- Miller, P. J. (1978). Forecasting with econometric methods: A comment. *Journal of Business*, **51**, 579–586.
- Muellbauer, J. N. J. (1994). The assessment: Consumer expenditure. *Oxford Review of Economic Policy*, **10**, 1–41.
- Newbold, P. (1993). Comment on ‘On the limitations of comparing mean squared forecast errors’, by M.P. Clements and D.F. Hendry. *Journal of Forecasting*, **12**, 658–660.
- Runkle, D. E. (1987). Vector autoregressions and reality. *Journal of Business and Economic Statistics*, **5**, 437–442.
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, **48**, 1–48. Reprinted in Granger, C. W. J. (ed.) (1990), *Modelling Economic Series*. Oxford: Clarendon Press.
- Sims, C. A., Stock, J. H., and Watson, M. W. (1990). Inference in linear time series models with some unit roots. *Econometrica*, **58**, 113–144.
- Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, **39**, 195–214.
- Taylor, J. B. (2000). The monetary transmission mechanism and the evaluation of monetary policy rules. Forthcoming, *Oxford Review of Economic Policy*.
- Turner, D. S. (1990). The role of judgement in macroeconomic forecasting. *Journal of Forecasting*, **9**, 315–345.
- White, H. (1990). A consistent model selection. In Granger, C. W. J. (ed.), *Modelling Economic Series*, pp. 369–383. Oxford: Clarendon Press.