CONSTRUCTIVE DATA MINING:
MODELING AUSTRALIAN INFLATION

Neil R. Ericsson*

Under revision: June 14, 2009

Abstract: This paper improves upon de Brouwer and Ericsson's (1998) model of Australian inflation in three directions: cointegration analysis, treatment of weak exogeneity, and model design. On the third issue, recent developments in computer-automated model selection help obtain a more parsimonious, empirically constant, data-coherent, encompassing error correction model for inflation in Australia. The level of consumer prices is a mark-up over domestic and import costs, with adjustments for dynamics and relative aggregate demand.

Keywords: Australia, cointegration, CPI, data mining, dynamic specification, encompassing, error correction, exogeneity, general to specific modeling, inflation, prices, model selection vector autoregression.

JEL classifications: C51, E31.

Neil R. Ericsson
Stop 15, Division of International Finance, Federal Reserve Board
2000 C Street, N.W., Washington, D.C. 20551 U.S.A.
(1) (202) 452–3709 (office)
(1) (202) 263–4850 (fax)
ericsson@frb.gov (email)
www.federalreserve.gov/research/staff/ericssonneilr.htm (webpage)

*The author is a staff economist in the Division of International Finance, Board of Governors of the Federal Reserve System, Washington, D.C. 20551 U.S.A. He may be reached on the Internet at ericsson@frb.gov. The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. I am grateful to Gordon de Brouwer for an earlier collaboration on this topic; see de Brouwer and Ericsson (1998). All numerical results were obtained using GiveWin Version 2.10, PcGive Version 10.1, and PcGets Version 1.02: see Doornik and Hendry (2001a, 2001b), Hendry and Doornik (2001), and Hendry and Krolzig (2001).

1 Introduction

To understand better the behavior of inflation and the role that a central bank may play in its determination, de Brouwer and Ericsson (1998) develop an empirical model of the Australian consumer price index (CPI) over 1977–1993. Their underlying economic theory is a mark-up model for prices, but the resulting empirical model also has elements relating to purchasing power parity and the Phillips curve. Their empirical model clarifies the relative importance of factors determining consumer price inflation. Further, the structure of the inflationary process in Australia does not appear to have changed over the 1980s and 1990s. Rather, the fall in inflation during the latter half of the 1980s and into the 1990s is explained in terms of changes in the determinants of inflation itself.

The model in de Brouwer and Ericsson (1998) does have two notable shortcomings. First, in their single-equation modeling of Australian inflation, de Brouwer and Ericsson (1998) augment the data from the underlying cointegration analysis with an impulse dummy (for a known measurement error) and the output gap (measured as detrended private final demand). While both variables are stationary in principle, their inclusion in the cointegration analysis could affect the results obtained, particularly because the output gap involves a linear trend. Second, an alternative single-equation model might have been obtained if a different model search path had been followed.

The current paper addresses both of these issues. First, cointegration is reanalyzed on the complete dataset, including the impulse dummy, private final demand, and a trend. The cointegrating vector in this expanded framework is similar to the one obtained by de Brouwer and Ericsson (1998), but unit labor costs are no longer weakly exogenous. That lack of weak exogeneity motivates an alternative single-equation modeling strategy that uses the Johansen estimates of the cointegrating vector when developing a single-equation model of inflation. Second, path dependence in model selection is addressed by using a recent multi-path model selection algorithm.

The implications are both methodological and empirical. This paper provides insights into the functioning of computer-automated model selection. By employing that tool, this paper also makes substantive improvements to a previously well-designed empirical model of inflation. The details of the model improvement highlight both the strengths and the limitations of computer-automated model selection.

This paper is organized as follows. Section 2 briefly describes the economic theory and the data. Section 3 summarizes the error correction model for Australian inflation that de Brouwer and Ericsson (1998) developed on quarterly series over 1977–1993. Section 4 re-analyzes the long-run properties of the underlying Australian CPI on an expanded dataset. Section 5 then designs a single-equation model of inflation using a recently developed algorithm for computer-automated model selection, as implemented in PcGets; see Hendry and Krolzig (2001). Depending upon the modeling strategy, the pre-search testing, the choice of required regressors, and the representation of the initial general model, PcGets obtains nine distinct—albeit similar—final specific models in its general-to-specific selection process. Additional analysis of those nine models results in a single final specification that encompasses and is even more parsimonious than the model in de Brouwer and Ericsson (1998). Section 6 shows that that final model is well-specified with empirically constant coefficients; and its economic interpretation is straightforward. Section 7 concludes.

2 Economic Theory and The Data

This section first discusses the mark-up model of price determination (Section 2.1) and then considers the data themselves (Section 2.2).

2.1 Economic Theory

While there are numerous theories and models of inflation, the enduring representation of the inflation process in Australia has been the mark-up model; see, for example, Richards and Stevens (1987) and Dwyer and Leong (2001). The mark-up model has a long-standing and continuing presence in economics generally; see Duesenberry (1950) and Franz and Gordon (1993) *inter alia*. The mark-up model is used throughout this paper, and it is general enough to embed several other well-known models, as noted below. This subsection describes the mark-up model, which underlies both this paper's empirical analysis and that in de Brouwer and Ericsson (1998).

In the long run, the domestic general price level is a mark-up over total unit costs, including unit labor costs, import prices, and energy prices. Assuming linear homogeneity, the long-run relation of the domestic consumer price level to its determinants is:

$$P = \mu \cdot (ULC^{\gamma})(IP^{\delta})(PET^{\kappa}). \tag{1}$$

The data are the underlying consumer price index (P), an index of the nominal cost of labor per unit of output (ULC), an index of tariff-adjusted import prices in domestic currency (IP), and an index of petrol prices in domestic currency (PET). The elasticities of the consumer price index with respect to ULC, IP, and PET are γ, δ, and κ, respectively, each of which is hypothesized to be greater than or equal to zero. The value $\mu - 1$ is the retail mark-up over costs, and both the mark-up and costs may vary over the cycle.

In practice, (1) is expressed in its log-linear form:

$$p = \ln(\mu) + \gamma \cdot ulc + \delta \cdot ip + \kappa \cdot pet, \tag{2}$$

where logarithms of variables are denoted by lower case letters. The log-linear form is used in the error correction model below. Linear homogeneity implies the following testable hypothesis:

$$\gamma + \delta + \kappa = 1, \tag{3}$$

which is unit homogeneity in all prices. Under that hypothesis, (2) can be rewritten as:

$$0 = \ln(\mu) + \gamma(ulc - p) + \delta(ip - p) + \kappa(pet - p), \tag{4}$$

which links real prices in the labor, foreign goods, and energy markets. This representation is particularly useful when interpreting empirical error correction models in the context of multiple markets influencing prices; cf. Juselius (1992), Durevall (1998), and Metin (1998). Additionally, through the term $(ip - p)$, equation (4) clarifies how the hypothesis of purchasing power parity is embedded in the mark-up model in (1). The empirical implementation also has ties to the Phillips curve by allowing the mark-up $\mu - 1$ to depend upon the output gap.

2.2 The Data

This section describes the data available and considers some of their basic properties. All data are quarterly, spanning 1976(3)–1993(3). Allowing for lags and transformations, estimation is over 1977(3)–1993(3) unless otherwise noted. de Brouwer and Ericsson (1995, Appendix 1) discuss in detail the definition and construction of the data.

The consumer price index is the central series of this study, and choosing an appropriate measure for it is complicated. The most publicly visible measure is the headline CPI (denoted P^h), published by the Australian Bureau of Statistics. However, the headline CPI includes a number of components that are subject to strong transitory fluctuations, that are controlled or influenced by the official sector, or that are unambiguously determined outside the Australian economy. While these components affect the Australian consumer, they are not necessarily readily modeled. While no final judgment exists as to which components should be excluded, this paper models one commonly used "underlying CPI" series, which is adjusted for such components. This underlying CPI (denoted P) is calculated as the headline CPI net of fresh fruit and vegetables, mortgage interest and consumer credit charges, automotive fuel, and health services. See the Reserve Bank of Australia (1994b) for a discussion of related issues and various measurements. In this paper, "CPI" always means this underlying CPI unless explicitly noted otherwise.

Figure 1 plots the quarterly inflation rates for underlying and headline CPI, denoted Δp and Δp^h.[1] The most noticeable differences between the two series are in

[1] The difference operator Δ is defined as $(1 - L)$, where the lag operator L shifts a variable one period into the past. Hence, for x_t (a variable x at time t), $Lx_t = x_{t-1}$ and so $\Delta x_t = x_t - x_{t-1}$.

1976 and 1984, when large changes in the cost of health services occurred. Figures 2a, 3a, and 4a plot the log of the CPI, its quarterly growth rate, and its annual growth rate respectively.

In the 1980s, the Australian inflation rate averaged around 8% per annum. At the beginning of the 1990s, the inflation rate fell substantially to average a much more moderate 2% per annum. This path of inflation has occurred alongside major changes in the Australian economy, including substantial reductions in tariffs, a shift to a more flexible and productivity-based system for setting wages, and a greater focus on international competitiveness. Moreover, the Reserve Bank of Australia has made a strong commitment to the preservation of low inflation, seeking to maintain an underlying annual inflation rate of around 2%–3%; see the Reserve Bank of Australia (1994a, p. 3).

Three additional series are of interest: ULC, IP, and PET. Figures 2a–c, 3a–c, and 4a–c plot the logs of these indices and their quarterly and annual growth rates, contrasting them with the corresponding transformations of the CPI. Over the sample as a whole, unit labor costs and import prices fall relative to the CPI, whereas petrol prices rise relative to the CPI. The time series for real petrol prices reflects the OPEC oil price increase in 1979, the fall in real oil prices in the latter half of the 1980s, and the dramatic but temporary increase in oil prices associated with the Gulf War.

One more series is of interest, private final demand (I). de Brouwer and Ericsson (1998) use this series to construct a proxy for the output gap, as measured by the residual i^{res} from regressing i on a constant and a trend t. Empirically, regression obtains:

$$i_t^{res} = i_t - 10.741 - 0.0069428t, \tag{5}$$

where the estimation sample is 1976(3)–1993(3) and a subscript t denotes time. Equation (5) implies that private final demand is growing at approximately 2.8% per annum on average. Section 4 includes both i and the linear trend directly in the cointegration analysis, whereas de Brouwer and Ericsson (1998) exclude both variables at that stage and add the output gap to the single-equation error correction model. As Section 4 shows, that choice has ramifications for the cointegration analysis itself.

The CPI and petrol prices are seasonally unadjusted, whereas unit labor costs, import prices, and private final demand are seasonally adjusted. While such "mixing" of data in the empirical analysis is unfortunate, the alternatives are limited, since the seasonally unadjusted CPI is the variable of interest, and unadjusted data are not available on all of the CPI's potential determinants. See Ericsson, Hendry, and Tran (1994) and the papers in Hylleberg (1992) for possible implications of using seasonally adjusted data in economic modeling.

More generally, $\Delta_j^i x_t = (1 - L^j)^i x_t$. If i (or j) is undefined, it is taken to be unity.

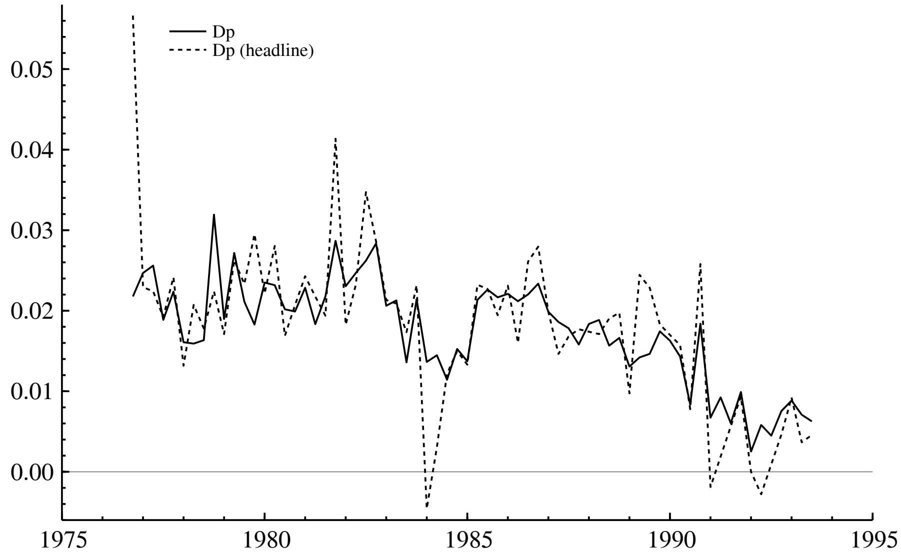

Figure 1: The underlying and headline CPI inflation rates.

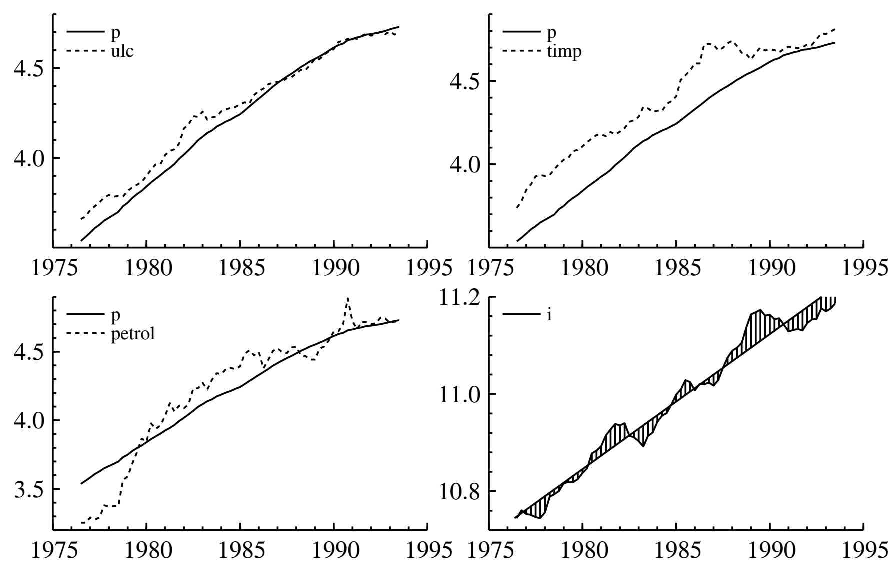

Figure 2: Log-levels of the basic series: p and ulc, p and ip, p and pet, and i.

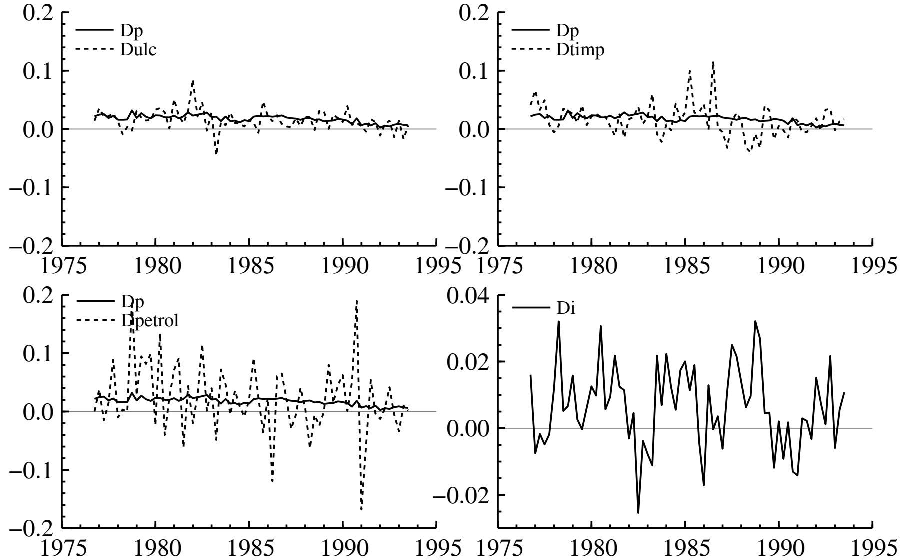

Figure 3: Quarterly growth rates of the basic series: Δp and Δulc, Δp and Δip, Δp and Δpet, and Δi.

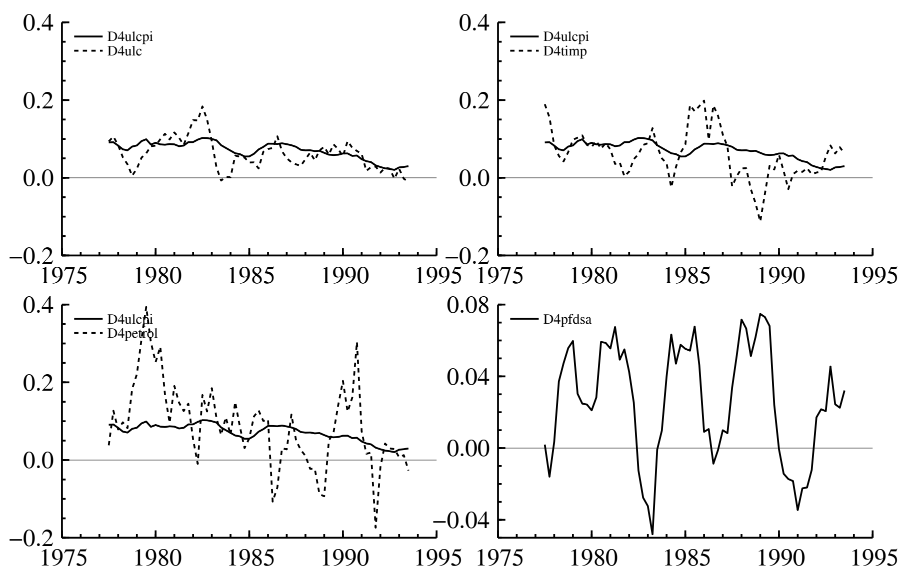

Figure 4: Annual growth rates of the basic series: $\Delta_4 p$ and $\Delta_4 ulc$, $\Delta_4 p$ and $\Delta_4 ip$, $\Delta_4 p$ and $\Delta_4 pet$, and $\Delta_4 i$.

3 Previous Results

This section summarizes the model of Australian inflation developed by de Brouwer and Ericsson (1998).

de Brouwer and Ericsson (1998) test for and find cointegration between the CPI, unit labor costs, import prices, and petrol prices; and the corresponding long-run relationship is homogeneous in prices. Unit labor costs, import prices, and petrol prices appear to be jointly weakly exogenous for the cointegrating vector, so de Brouwer and Ericsson (1998) model the CPI as a single-equation conditional error correction model (ECM), obtained from an autoregressive distributed lag for the CPI. The finally selected ECM in de Brouwer and Ericsson (1998) is:

$$\begin{aligned}
\Delta p_t \quad = \quad & + \underset{[0.0062]}{0.0141} \; \Delta pet_t + \underset{[0.0086]}{0.0763} \; i^{res}_{t-1} \\
& - \underset{[0.0066]}{0.0415} \; (p - ulc)_{t-1} - \underset{[0.0035]}{0.0394} \; (p - ip)_{t-1} \\
& - \underset{[0.0022]}{0.0082} \; (p - pet)_{t-1} + \underset{[0.0013]}{0.0096} \; D_t + \underset{[0.00058]}{0.00749} \\
& - \underset{[0.0009]}{0.0017} \; S_{1t} - \underset{[0.0011]}{0.0009} \; S_{2t} - \underset{[0.0009]}{0.0021} \; S_{3t}
\end{aligned} \tag{6}$$

$T = 65$ [1977Q3–1993Q3] $\quad R^2 = 0.87 \quad \hat{\sigma} = 0.251\%$

$DW = 1.98 \quad AR: F(5, 50) = 0.50$ [0.7052]

$ARCH: F(4, 47) = 0.85$ [0.4989] $\quad Normality: \chi(2) = 2.24$ [0.3260]

$HeteroA: F(14, 40) = 0.99$ [0.4834] $\quad RESET: F(1, 54) = 0.22$ [0.6422],

which is written in two different representations as their equations (10) and (14). Equation (6) is parsimonious and empirically constant, and its economic interpretation is straightforward because it solves directly for the long-run solution in equation (2). As an error correction model, equation (6) captures long-run effects that were ignored in some previous models of the Australian CPI, which were in first differences only. Including the error correction term in the empirical model of Australian CPI ties the model more closely to its theoretical underpinnings and improves the goodness-of-fit.

Equation (6) has been used for forecasting Australian inflation, and its simple structure (in essence, by depending on only lagged variables) makes one-step ahead forecasts easy. Also, contemporary models of Australian inflation developed at the Reserve Bank of Australia are very similar in structure to (6), even when using more recent data: see Dwyer and Leong (2001). In spite of the apparent robustness of (6) over time, its development had two shortcomings: the exclusion of private final demand from the cointegration analysis, and possible dependence on the path taken for model selection. The remainder of the current paper addresses these two issues to obtain an improved specification.

4 Integration and Cointegration

This section presents unit root tests for the variables of interest (Section 4.1). Then, Johansen's maximum likelihood procedure is applied to test for cointegration among the CPI, unit labor costs, import prices, petrol prices, private final demand, and a linear trend (Section 4.2). Long-run price homogeneity and the adjustment mechanism are examined in the Johansen framework.

4.1 Integration

Table 1 lists ADF statistics and related calculations for the data. In order to test whether a given series is I(0), I(1), I(2), or I(3), Table 1 calculates unit root tests for the original variables (all in logs), for their changes, and for the changes of the changes. This permits testing the order of integration, albeit by testing adjacent orders of integration in a pairwise fashion. The largest estimated root ($\hat{\rho}$) appears adjacent to each Dickey-Fuller statistic: this root should be approximately unity if the null hypothesis is correct. The lag length of the ADF regression is based on minimizing the AIC, starting with a maximum of four lags.

Empirically, all variables appear to be integrated of order two or lower. If inferences are made on the reported Dickey-Fuller statistics alone, then unit labor costs, import prices, petrol prices, and private final demand appear to be I(1), whereas the CPI appears to be I(2). However, the estimated root for Δp is only 0.650, which numerically is much less than unity. Also, while the *first*-order ADF minimizes the AIC for Δp, the *zeroth*-order ADF for Δp has nearly the same value of the AIC (-11.09 versus -11.17), but the value of the ADF(0) is -4.46, significant at the 1% level. Thus, all four price series are treated below as if they are I(1), while recognizing that some caveats may apply.

4.2 Cointegration

Cointegration analysis helps clarify the long-run relationships between integrated variables. Johansen's (1988, 1991) procedure is maximum likelihood for finite-order vector autoregressions (VARs) and is easily calculated for such systems, so it is used here. Empirically, the lag order of the VAR is not known *a priori*, so some testing of lag order may be fruitful in order to ensure reasonable power of the Johansen procedure. Given the number of variables and the number of observations involved, the largest system considered is a fourth-order VAR of p, ulc, ip, pet, and i. In that VAR, a linear trend is restricted to lie in the cointegration space, and an intercept and seasonal dummies enter freely. Table 2 shows that it is statistically acceptable to simplify that fourth-order VAR to a first-order VAR.

Table 3 reports the standard statistics and estimates for Johansen's procedure ap-

Table 1: ADF statistics for testing a unit root in various time series.

Variable x	lag ℓ	$t_{ADF(\ell)}$	$\hat{\rho}$	$\hat{\sigma}$ (%)	t-prob (%)	F-prob (%)	AIC
				Null hypothesis: x_t is I(1)			
p	2	0.46	1.007	0.359	4.2	34.0	–11.14
ulc	0	–0.38	0.982	1.785	—	60.8	–7.96
ip	0	–1.15	0.952	2.685	—	67.5	–7.14
pet	1	–2.30	0.880	5.626	18.1	79.0	–5.65
i	3	–3.12	0.816	1.137	15.1	81.0	–8.82
ecm	1	–2.27	0.879	1.692	14.2	99.2	–8.05
				Null hypothesis: x_t is I(2)			
p	1	–2.90	0.650	0.356	1.5	51.9	–11.17
ulc	0	–7.26**	0.039	1.786	—	30.9	–7.96
ip	0	–6.77**	0.107	2.701	—	24.3	–7.13
pet	0	–9.28**	–0.208	5.837	—	73.8	–5.59
i	0	–5.65**	0.271	1.212	—	54.7	–8.73
ecm	4	–2.37	0.330	1704	1.4	—	–8.00
				Null hypothesis: x_t is I(2)			
p	0	–12.78**	–0.485	0.379	—	26.9	–11.06
ulc	4	–5.35**	–1.951	1.915	3.5	—	–7.76
ip	3	–7.62**	–2.023	2.797	0.0	54.5	–7.02
pet	4	–6.59**	–3.408	6.448	1.7	—	–5.34
i	1	–8.17**	–0.797	1.335	5.5	42.2	–8.53
ecm	3	–8.70**	–2.361	1.777	0.0	90.2	–7.93

Notes:

1. Fourth-order ADF regressions were initially estimated, and the final lag length was selected to minimize the Akaike Information Criterion (AIC). The columns report the name of the variable examined, the selected lag length ℓ, the ADF statistic on the simplified regression ($t_{ADF(\ell)}$), the estimated coefficient on the lagged level that is being tested for a unit value ($\hat{\rho}$), the regression's residual standard error ($\hat{\sigma}$), the tail probability of the t-statistic on the longest lag of the final regression (t-prob, in %), the tail probability of the F-statistic for the lags dropped (F-prob, in %), and the AIC.

2. All of the ADF regressions include both an intercept and a linear trend. MacKinnon's (1991) approximate finite-sample critical values for the corresponding ADF statistics are –3.17 (10%), –3.48 (5%), and –4.11 (1%) for $T = 62$. Rejection of the indicated null hypothesis is denoted by $^+$, *, and ** for the 10%, 5%, and 1% levels.

Table 2: F and related statistics for the sequential reduction from the fourth-order VAR to the first-order VAR.

Null hypothesis						Maintained hypothesis		
System	k	$\log(\mathcal{L})$	SC	HQ	AIC	VAR(4)	VAR(3)	VAR(2)
VAR(4)	130	1008.28	–22.68	–25.31	–27.02	—		
						—		
↓						—		
VAR(3)	105	985.77	–23.59	–25.72	–27.10	$F(25, 131)$		
						1.08		
↓						[0.376]		
VAR(2)	80	963.81	–24.52	–26.14	–27.19	$F(50, 162)$	$F(25, 150)$	
						1.140	1.20	
↓						[0.268]	[0.250]	
VAR(1)	55	951.21	–25.74	–26.85	–27.58	$F(75, 171)$	$F(50, 185)$	$F(25, 168)$
						1.01	0.98	0.74
						[0.460]	[0.526]	[0.808]

Notes:

1. The first six columns report the VAR with its order and, for each system, the number of unrestricted parameters k, the log-likelihood $\log(\mathcal{L})$, the Schwarz criterion SC, the Hannan-Quinn criterion HQ, and the Akaike Information Criterion AIC.

2. In the three last columns, the three entries within a given block of numbers are the approximate F-statistic for testing the null hypothesis (indicated by the model to the far left of the entry) against the maintained hypothesis (indicated by the model above the entry), the value of that F-statistic, and the tail probability associated with that value of the F-statistic (in brackets).

plied to this first-order VAR. The maximal eigenvalue and trace eigenvalue statistics (λ_{max} and λ_{trace}) strongly reject the null of no cointegration in favor of at least one cointegrating relationship, and little evidence exists for more than one. Parallel statistics with a degrees-of-freedom adjustment (λ^a_{max} and λ^a_{trace}) give a similar picture, reflecting one very large eigenvalue (0.854) and four much smaller eigenvalues.

Table 3 also reports the standardized eigenvectors and adjustment coefficients, denoted β' and α in a common notation. The first row of β' is the estimated cointegrating vector, which can be written in the form of (2):

$$\begin{aligned} p \quad = \quad & \ln(\widehat{\mu}) \; + \; 0.411\,ulc \; + \; 0.507\,ip \; + \; 0.101\,pet \\ & - \; 0.893\,i + 0.00685t. \end{aligned} \tag{7}$$

All coefficients have their anticipated signs. Numerically, the coefficients on ulc and ip are approximately equal in value, and the sum of price coefficients in (7) is numerically close to unity, i.e., 1.019. Statistically, the restriction of long-run unit price homogeneity cannot be rejected: $\chi^2(1) = 0.02$ [0.884]; see Johansen and Juselius (1990), Johansen (1995), and Doornik and Hendry (2001b) for the form of the test. The asymptotic null distribution is denoted by $\chi^2(\cdot)$ with degrees of freedom in parentheses, and the asymptotic p-value is in square brackets.

Also, in (7), the ratio of the coefficient on the trend to that on private final demand is virtually the slope coefficient in equation (5): -0.00767 versus -0.00694. That is, the numerical role of the trend in the cointegrating relationship (7) appears to be solely to remove the trend from private final demand, and not from any other variables. That result is consistent with ECM estimates in de Brouwer and Ericsson (1998), where the trend enters only through the output gap.

In Table 3, the coefficients in the first column of α measure the feedback effect of the (lagged) disequilibrium in the cointegrating relation onto the variables in the vector autoregression. Specifically, -0.079 is the estimated feedback coefficient for the CPI equation. The negative coefficient implies that an "excess" mark-up induces a lower CPI inflation rate. The coefficient's numerical value entails gradual adjustment to remaining disequilibrium and so substantial smoothing of unit labor costs, import prices, and petrol prices in obtaining the CPI.

The next block in Table 3 reports values of the statistic for testing weak exogeneity of a given variable for the cointegrating vector. Equivalently, the statistic tests whether or not the corresponding row of α is zero; see Johansen (1995). If it is zero, disequilibrium in the cointegrating relationship does not feed back onto that variable. Individually and together, import prices, petrol prices, and private final demand are weakly exogenous: $\chi^2(3) = 2.01$ [0.569]. Imposing weak exogeneity of those variables jointly with long-run homogeneity also is not rejected: $\chi^2(4) = 2.05$ [0.727]. However, unit labor costs are not weakly exogenous, contrary to what de Brouwer and Ericsson (1998) found in a VAR with just p, ulc, ip, and pet.

Table 3: A cointegration analysis of the Australian price data.

rank of π	$r=0$	$r\leq 1$	$r\leq 2$	$r\leq 3$	$r\leq 4$	$r\leq 5$
$\log(\mathcal{L}_r)$	863.29	925.74	938.88	945.99	949.23	951.21
Eigenvalue λ_r	–	0.854	0.333	0.196	0.095	0.059

	Null hypothesis				
	$r=0$	$r\leq 1$	$r\leq 2$	$r\leq 3$	$r\leq 4$
$\lambda_{\max}$	124.9**	26.29	14.21	6.47	3.97
$\lambda^a_{\max}$	115.3**	24.27	13.12	5.98	3.66
95% critical value	37.52	31.46	25.54	18.96	12.25
λ_{trace}	175.8**	50.95	24.66	10.44	3.97
λ^a_{trace}	162.3**	47.03	22.76	9.64	3.66
95% critical value	87.31	62.99	42.44	25.32	12.25

	Eigenvectors β'					
Variable	*p*	*ulc*	*ip*	*pet*	*i*	*trend*
	1	–0.411	–0.507	–0.101	–0.893	0.00685
	–1.207	1	0.238	–0.038	0.036	0.00334
	2.946	–1.167	1	–0.128	–6.946	0.00008
	–1.978	–0.817	0.020	1	0.028	0.02960
	0.316	–0.281	0.727	–0.043	1	–0.01529

Variable	Adjustment coefficients α				
p	–0.079	0.010	–0.002	–0.001	–0.001
ulc	–0.091	–0.238	–0.026	0.002	0.010
ip	–0.050	0.108	–0.015	0.031	–0.060
pet	–0.102	–0.205	–0.040	–0.122	–0.077
i	–0.002	–0.190	0.020	0.003	–0.003

	Weak exogeneity test statistics				
	p	*ulc*	*ip*	*pet*	*i*
$\chi^2(1)$	96.6**	8.42**	1.14	1.04	0.01

	Statistics for testing the significance of a given variable in $\beta' x$					
	p	*ulc*	*ip*	*pet*	*i*	*trend*
$\chi^2(1)$	11.0**	4.65*	32.7**	7.28**	24.7**	9.14**

	Multivariate statistics for testing trend stationarity				
	p	*ulc*	*ip*	*pet*	*i*
$\chi^2(4)$	51.4**	62.9**	31.5**	54.7**	48.0**

The penultimate block in Table 3 reports chi-squared statistics for testing the significance of individual variables in the cointegrating vector. Each variable is significant, including the linear trend.

The final block in Table 3 reports values of a multivariate statistic for testing the trend stationarity of a given variable. Specifically, this statistic tests the restriction that the cointegrating vector contains all zeros except for a unity corresponding to the designated variable and an unrestricted estimate on the trend, with the test being conditional on the presence of exactly one cointegrating vector; see Johansen (1995, p. 74). For instance, the null hypothesis of a trend-stationary CPI implies that the cointegrating vector is $(1\ 0\ 0\ 0\ 0\ *)'$, where "$*$" represents an unrestricted coefficient on the linear trend. Empirically, all of the stationarity tests reject with p-values less than 0.01%. By being multivariate, these statistics may have higher finite sample power than their univariate counterparts. Also, the null hypothesis is the stationarity of a given variable rather than the nonstationarity thereof, and stationarity may be a more appealing null hypothesis. That said, these rejections of stationarity are in line with the *inability* in Table 1 to reject the null hypothesis of a unit root in each of p, ulc, ip, pet, and i.

Table 4 reports the estimated values of α and β when estimated unrestrictedly and when estimated with homogeneity imposed, with weak exogeneity imposed, and with both homogeneity and weak exogeneity imposed. Not only are these hypotheses statistically acceptable, the estimates with both homogeneity and weak exogeneity imposed are virtually unchanged *numerically* from the unrestricted estimates or from those obtained by imposing subsets of the hypotheses. The similarity of coefficient estimates across the various restrictions points to the robustness of the results and is partial evidence in favor of those restrictions.

Because neither the CPI nor unit labor costs are weakly exogenous for the cointegrating vector, inferences in a single equation of the CPI are hazardous if the cointegrating vector is estimated jointly with the equation's dynamics; see Hendry (1995). One solution is to model the CPI and unit labor costs as a subsystem, conditional upon import prices, petrol prices, and private final demand. A second solution—adopted below—is to construct an error correction term from the system estimates and then to develop a single-equation error correction model that uses that system-based error correction term. This second approach is adopted below. See Hendry and Doornik (1994) and Juselius (1992) for paradigms of these two approaches.

5 Computer-automated Model Selection

This section first describes the model selection algorithm in PcGets (Section 5.1) and then applies that algorithm to an ECM representation of an autoregressive distributed lag model for Australian inflation (Section 5.2). That ECM representation—and

Table 4: Just-identified and over-identified estimates of β and α, with corresponding estimated standard errors, from a cointegration analysis of Australian CPI.

Variable corresponding to an element of β' or α'					
p	*ulc*	*ip*	*pet*	*i*	*trend*
Estimate of β' (just-identified)					
1	−0.411	−0.507	−0.101	−0.893	0.00685
	(0.106)	(0.052)	(0.036)	(0.139)	(0.00150)
Estimate of β' (long-run homogeneity imposed)					
1	−0.399	−0.499	−0.102	−0.888	0.00648
	(0.052)	(0.047)	(0.031)	(0.130)	(0.00093)
Estimate of β' (weak exogeneity imposed)					
1	−0.415	−0.508	−0.100	−0.880	0.00676
	(0.104)	(0.051)	(0.035)	(0.136)	(0.00147)
Estimate of β' (long-run homogeneity and weak exogeneity imposed)					
1	−0.402	−0.497	−0.101	−0.872	0.00630
	(0.051)	(0.046)	(0.030)	(0.128)	(0.00091)
Estimate of α' (just-identified)					
−0.079	−0.091	−0.050	−0.102	−0.002	
(0.005)	(0.031)	(0.049)	(0.104)	(0.023)	
Estimate of α' (long-run homogeneity imposed)					
−0.081	−0.093	−0.051	−0.103	−0.002	
(0.005)	(0.031)	(0.050)	(0.105)	(0.023)	
Estimate of α' (weak exogeneity imposed)					
−0.079	−0.090	0	0	0	
(0.004)	(0.030)				
Estimate of α' (long-run homogeneity and weak exogeneity imposed)					
−0.081	−0.093	0	0	0	
(0.004)	(0.031)				

all of its simplifications—use the most restricted error correction term in Table 4, i.e., imposing both homogeneity and the weak exogeneity of import prices, petrol prices, and private final demand. For convenience below, that error correction term is denoted *ecm*.

5.1 PcGets's Algorithm

Hendry and Krolzig (2001) develop a computer program PcGets, which extends and improves upon Hoover and Perez's (1999) automated model-selection algorithm; see also Hendry and Krolzig (1999, 2003a, 2003b) and Krolzig and Hendry (2001). PcGets utilizes one-step and multi-step simplifications along multiple paths, diagnostic tests as additional checks on the simplified models, and encompassing tests to resolve multiple terminal models. Both analytical and Monte Carlo evidence show that the resulting model selection is relatively non-distortionary for Type I errors. At an intuitive level, PcGets functions as a series of sieves that aims to retain parsimonious congruent models while discarding both noncongruent models and over-parameterized congruent models. This feature of the algorithm is eminently sensible, noting that the DGP itself is congruent and is as parsimonious as is feasible.

The current subsection summarizes PcGets as an automated model-selection algorithm, thereby providing the necessary background for interpreting its application in Section 5.2. For ease of reference, the algorithm in PcGets is divided into four "stages", denoted Stage 0, Stage 1, Stage 2, and Stage 3. For full details of PcGets's algorithm, see Hendry and Krolzig (2001, Appendix A1). Hendry and Krolzig (2002) describe the relationship of the general-to-specific approach to other modeling approaches in the literature, and Hoover and Perez (2000) extend the general-to-specific approach to cross-section regressions.

Stage 0: the general model and F pre-search tests. Stage 0 involves two parts: the estimation and evaluation of the general model, and some pre-search tests aimed at simplifying the general model before instigating formal multi-path searches.

First, the general model is estimated, and diagnostic statistics are calculated for it. If any of those diagnostic statistics is unsatisfactory, the modeler must decide what to do next—whether to "go back to the drawing board" and develop another general model, or whether to continue with the simplification procedure, perhaps ignoring the offending diagnostic statistic or statistics.

Second, PcGets attempts to drop various sets of potentially insignificant variables. PcGets does so by dropping all variables at a given lag, starting with the longest lag. PcGets also does so by ordering the variables by the magnitude of their t-ratios and either dropping a group of individually insignificant variables or (alternatively) retaining a group of statistically significant variables. In effect, an F pre-search test for a group of variables is a single test for multiple simplification paths, a characteristic that helps control the costs of search. If these tests result in a statistically satisfactory reduction of the general model, then that new model is the starting point for Stage 1. Otherwise, the general model itself is the starting point for Stage 1.

Stage 1: a multi-path encompassing search. Stage 1 tries to simplify the model from Stage 0 by employing multi-path searches, all the while ensuring that the diagnostic tests are not rejected. If all variables are individually statistically significant, then the initial model in Stage 1 is the final model. If some variables are statistically insignificant, then PcGets tries deleting those variables to obtain a simpler model. PcGets proceeds down a given simplification path only if the models along that path have satisfactory diagnostic statistics. If a simplification is rejected or if a diagnostic statistic fails, PcGets backtracks along that simplification path to the most recent previous acceptable model and then tries a different simplification path. A terminal model results if the model's diagnostic statistics are satisfactory and if no remaining regressors can be deleted.

If PcGets obtains only one terminal model, then that model is the final model, and PcGets proceeds to Stage 3. However, because PcGets may pursue multiple simplification paths in Stage 1, PcGets may obtain multiple terminal models. To resolve such a situation, PcGets creates a union model from those terminal models and tests each terminal model against that union model. PcGets then creates a new union model, which nests all of the surviving terminal models; and that union model is passed on to Stage 2.

Stage 2: another multi-path encompassing search. Stage 2 in effect reruns Stage 1, applying the simplification procedures from Stage 1 to the union model obtained from Stage 1. The resulting model is the final model. If Stage 2 obtains more than one terminal model after applying encompassing tests, then the final model is selected by using the Akaike, Schwarz, and Hannan-Quinn information criteria; see Akaike (1973, 1981), Schwarz (1978), and Hannan and Quinn (1979) for the design of these information criteria, and Atkinson (1981) for the relationships between them.

Stage 3: subsample evaluation. Stage 3 re-estimates the final model over two subsamples and reports the results. If a variable is statistically significant in the full sample and in both subsamples, then the inclusion of that variable in the final model is regarded as "100% reliable". If a variable is statistically insignificant in one or both subsamples or in the full sample, then its measure of reliability is reduced. A variable that is statistically insignificant in both subsamples and in the full sample is regarded as being "0% reliable". The modeler is left to decide what action, if any, to take in light of the degree of reliability assigned to each of the regressors.

PcGets thus has two components:

1. Estimation and diagnostic testing of the general unrestricted model (Stage 0); and

2. Selection of the final model by
   (a) pre-search simplification of the general unrestricted model (Stage 0),
   (b) multi-path (and possibly iterative) selection of the final model (Stages 1 and 2), and
   (c) post-search evaluation of the final model (Stage 3).

This subsection's description of these four stages summarizes the algorithm in PcGets. Below, Section 5.2 summarizes the actual simplification paths taken by PcGets in practice, thereby providing additional insight into PcGets's algorithm.

PcGets requires the modeler to choose which tests are calculated and to specify the critical values for those tests. In PcGets, the modeler can choose the test statistics and their critical values directly, although doing so is tedious because of the number of statistics involved. To simplify matters, PcGets offers two options with pre-designated selections of test statistics and critical values. These two options are called "liberal" and "conservative" model selection strategies. The liberal strategy errs on the side of keeping some variables, even although they may not actually matter. The conservative strategy keeps only variables that are clearly significant statistically, erring in the direction of excluding some variables, even although those variables may matter. When modeling economic data, the liberal strategy typically appears preferable, although (as below) the two approaches may generate similar or identical results on specific datasets.

5.2 Modeling of Australian Inflation Revisited

Using PcGets, the current subsection assesses the possible path dependence of equation (6). Four choices in the model selection process allow further investigation of that equation's robustness and of PcGets's algorithm itself. Those choices concern the following.

1. Model strategy: either liberal ("L") or conservative ("C").
2. Pre-search testing (Stage 0): either switched on ("Yes") or off ("No").
3. The choice of required (a.k.a. "fixed") regressors: either a fixed intercept and seasonals ("Yes"), or no fixed regressors ("No"). This facility allows the user to force PcGets to always include certain variables.
4. The representation of the initial general ECM: either with Δp_t, or with $\Delta_4 p_t$, as explained below.

This subsection estimates the initial general model and allows PcGets to simplify that model under each of the 16 permutations implied by this list of choices. PcGets obtains nine distinct models; and additional analysis results in a single final specification that encompasses and is more parsimonious than the model in de Brouwer and Ericsson (1998). Section 6 shows that that final model is well-specified with empirically constant coefficients; and its economic interpretation is straightforward.

Table 5 lists the estimates and standard errors for the ECM representation of the fourth-order ADL model of the CPI, unit labor costs, import prices, petrol prices, and private final demand. The standard diagnostic statistics do not reject. The coefficient on the error correction term ecm_{t-1} appears to be highly significant statistically, with a t-ratio of -4.09. Because Table 3 shows that the four prices and private final demand are cointegrated with a disequilibrium error measured by the variable ecm, the t-ratio on ecm_{t-1} is (asymptotically) distributed as a standardized normal variate, although with a nonzero mean if long-run feedback occurs in the inflation equation. The magnitude of that t-ratio indicates that such feedback is present.

As mentioned above, the representation of the initial general ECM can affect the model selected by PcGets. In its simplification process, PcGets imposes only "zero restrictions", i.e., PcGets can set coefficients to be equal to only zero. Although a linear model is invariant to nonsingular linear transformations of its data, the coefficients of that model are *not* invariant under such transformations. For example, a given model is invariant to including the regressors x_t and x_{t-1} or the regressors Δx_t and x_{t-1}, but the deletion of x_{t-1} results in two different simplifications.

For Table 5, two representations are of particular interest: one with Δi_t, and the other with $\Delta_4 i_t$. As background, de Brouwer and Ericsson (1998) begin their simplification path from an ECM almost identical to Table 5, but with homogeneity imposed on the long-run price coefficients and with the restriction in equation (5) imposed on the coefficients of i and the trend. de Brouwer and Ericsson (1998) then simplify their unrestricted ECM to the parsimonious ECM in equation (6). That said, de Brouwer and Ericsson (1995, Appendix 3) note an alternative parsimonious ECM that results from simplifying an equivalent unrestricted ECM. Their alternative ECM includes the term $\Delta_4 i_t$, which (aside from a constant term) is interpretable as the change in the output gap in (5).

For the model in Table 5, the transformation from Δi_t to $\Delta_4 i_t$ proceeds as follows. The variable $\Delta_4 i_t$ is the sum of the variables Δi_t, Δi_{t-1}, Δi_{t-2}, and Δi_{t-3}, all of which appear in Table 5. Thus, the four variables $\{\Delta i_t, \Delta i_{t-1}, \Delta i_{t-2}, \Delta i_{t-3}\}$ in Table 5 can be transformed to $\{\Delta_4 i_t, \Delta i_{t-1}, \Delta i_{t-2}, \Delta i_{t-3}\}$ without altering the model itself. However, PcGets may obtain different models, depending upon which way in which the dynamics of i are represented. Numerically, $\{\Delta i_t, \Delta i_{t-1}, \Delta i_{t-2}, \Delta i_{t-3}\}$ enter the model in Table 5 as:

$$-0.046\,\Delta i_t - \ 0.050\,\Delta i_{t-1} - \ 0.012\,\Delta i_{t-2} - \ 0.024\,\Delta i_{t-3}, \tag{8}$$

Table 5: An unrestricted ECM representation for the fourth-order autoregressive distributed lag model of Australian inflation, conditional on unit labor costs, import prices, petrol prices, and real private final demand.

Variable	Lag j			
	0	1	2	3
Δp_{t-j}	-1 $(-)$	-0.057 (0.132)	0.125 (0.124)	-0.085 (0.129)
Δulc_{t-j}	0.030 (0.024)	-0.019 (0.025)	-0.015 (0.025)	-0.004 (0.025)
Δip_{t-j}	0.016 (0.014)	-0.014 (0.016)	-0.008 (0.017)	-0.003 (0.015)
Δpet_{t-j}	0.012 (0.007)	0.002 (0.006)	-0.002 (0.006)	0.003 (0.006)
Δi_{t-j}	-0.046 (0.036)	-0.050 (0.045)	-0.012 (0.040)	-0.024 (0.038)
ecm_{t-j}		-0.087 (0.021)		
D_{t-j}	0.010 (0.003)			
intercept, S_{jt}	-0.784 (0.192)	-0.0019 (0.0012)	-0.0010 (0.0011)	-0.0021 (0.0012)

$T = 65$ [1977Q3–1993Q3] $RSS = 0.0002653$ $R^2 = 0.899$ $\hat{\sigma} = 0.258\%$

$DW = 2.04$ $AR : F(5, 35) = 0.32$ [0.8968]

$ARCH : F(4, 32) = 0.44$ [0.7768] $Normality : \chi^2(2) = 0.64$ [0.7270]

$HeteroA : \chi^2(44) = 40.6$ [0.6194] $RESET : F(1, 39) = 0.00$ [0.9562] whereas the variables $\{\Delta_4 i_t, \Delta i_{t-1}, \Delta i_{t-2}, \Delta i_{t-3}\}$ would enter the model in Table 5 as:

$$-0.046\,\Delta_4 i_t - \ 0.005\,\Delta i_{t-1} + \ 0.033\,\Delta i_{t-2} + \ 0.022\,\Delta i_{t-3}. \tag{9}$$

A zero restriction on Δi_{t-1} in equation (9) might be statistically reasonable, whereas zero restriction on the same variable in equation (8) might not, noting that their respective coefficients are -0.005 and -0.050.

Table 6 summarizes PcGets's model simplifications under the 16 different scenarios described above. Several features of the simplifications are notable. First, the seasonal dummies are invariably dropped if they are not forced to enter the model (i.e., if they are not "fixed"). Thus, the model's seasonality appears to be purely "internal", in that the regression combines the right-hand side variables in such a way that their combined seasonality just offsets or matches the seasonality in the dependent variable. Second, pre-search testing markedly reduces the number of paths that need to be searched in Stage 1. In some scenarios, pre-search testing obtains the final model. As a consequence, pre-search testing frequently reduces or eliminates the occurrence of multiple terminal models. Third, model simplifications in all scenarios focus on a small set of variables: only 10 variables, compared to 25 variables in Table 5. Of those variables, three are always included in the final model: the intercept, the error correction term ecm_{t-1}, and the dummy variable D_t. The rate of change of petrol prices (Δpet_t) is almost always included in the final model, as is some form of output growth—either Δi_t or $\Delta_4 i_t$.

To sort out which of the nine models in Table 6 is to be preferred, Table 7 lists the equation standard error and information criteria for each model, along with the F-statistics of each model against two "union" models, denoted models M11 and M12. Model M11 is the minimal union of models M1–M9. Model M12 is the minimal union of models M1–M9 and equation (6). That equation is the final model from de Brouwer and Ericsson (1998), and it is denoted Model M10 on Table 7.

Models M1–M9 can be winnowed, as follows. Models M4, M5, M6, and M9 are each rejected at the 95% level relative to one of the two union models—M11 or M12—or both. Model M1 parsimoniously simplifies to model M2, which is model M1 without the seasonal dummies. Model M7 likewise simplifies to model M8.

Only three models remain: Models M2, M3, and M8. All three include, Δpet_t, ecm_{t-1}, the intercept, and the dummy variable D_t. They differ in that model M8 includes $\Delta_4 i_t$, whereas model M2 includes Δi_t, and model M3 includes Δi_t and Δp_{t-2}. There appears little to distinguish the three models statistically, a situation that parallels one observed by Boughton (1993) and Hendry and Starr (1993) on some competing narrow money demand equations for the United States. Model M8 is slightly simpler to interpret economically, and it is slightly more parsimonious than model M3, so Section 6 briefly examines model M8 as the "final" specification. Additional data and new tests may resolve such an empirical ambiguity.

Table 6: Computer-automated selection of models for Australian inflation, categorized according to model strategy, pre-search testing, fixed regressors, and representation of the general model.

Model strategy	Pre-search?	Fixed regressors?	k_1	k_f	Number of paths	Number of final models	Δpet_t	Δi_t	$\Delta_4 i_t$	Δp_{t-2}	ecm_{t-1}	S_0	D_t	S_{jt}	$\hat{\sigma}$ (%)	Model name
Model representation: Δi_t in the general model																
L	Yes	Yes	8	8	0	1	•	•			•	•	•	•	0.237	M1
L	No	Yes	25	8	26	5	•	•			•	•	•	•	0.237	M1
L	Yes	No	6	5	4	3	•	•			•	•	•		0.247	M2
L	No	No	25	6	29	6	•	•		•	•	•	•		0.239	M3
C	Yes	Yes	6	6	0	1					•	•	•	•	0.258	M4
C	No	Yes	25	7	26	2		•			•	•	•	•	0.245	M5
C	Yes	No	4	4	0	1	•				•	•	•		0.256	M6
C	No	No	25	4	29	2	•				•	•	•		0.256	M6
Model representation: $\Delta_4 i_t$ in the general model																
L	Yes	Yes	8	8	0	1	•		•		•	•	•	•	0.235	M7
L	No	Yes	25	8	27	5	•		•		•	•	•	•	0.235	M7
L	Yes	No	6	4	4	2	•				•	•	•		0.256	M6
L	No	No	25	5	30	5	•		•		•	•	•		0.243	M8
C	Yes	Yes	7	7	0	1			•		•	•	•	•	0.246	M9
C	No	Yes	25	7	27	2			•		•	•	•	•	0.246	M9
C	Yes	No	5	5	0	1	•		•		•	•	•		0.243	M8
C	No	No					•		•		•	•	•		0.243	M8

Notes:

1. The model strategy is either "liberal" (L) or "conservative" (C); pre-search testing is either included ("Yes") or excluded ("No"); and the intercept S_0 and the seasonal dummies S_{jt} are either fixed in the regression ("Yes") or allowed to be deleted in model selection ("No"). The number of regressors after pre-search testing is k_1 (i.e., at the beginning of Stage 1); and the number of regressors in the final specific model is k_f. The "number of paths" is the number of different simplification paths considered in Stage 1; and the "number of final models" is the number of distinct terminal specifications obtained in Stage 1.

2. A dot "•" under a variable denotes that that variable appears in the final specific model for that model strategy, pre-search and fixed-variables options, and representation of the general model. The equation standard error of that final specific model is $\hat{\sigma}$; and its model name is one of M1, ..., M9.

Table 7: Information criteria for the various final specific models in Table 6, and F-statistics for testing those models against two more general union models—models M11 and M12.

Model	Information criterion				Maintained hypothesis		
name	$\hat{\sigma}$	SC	HQ	AIC	Table 5	M11	M12
M1	0.237%	–8.869	–9.031	–9.137	[0.944]	$F(2,55)$ 1.74 [0.1853]	$F(6,51)$ 0.73 [0.6297]
M2	0.247%	–8.932	–9.033	–9.099	[0.753]	$F(5,55)$ 2.32 [0.0555]	$F(9,51)$ 1.34 [0.2410]
M3	0.239%	–8.947	**–9.068**	–9.148	[0.907]	$F(4,55)$ 1.63 [0.1798]	$F(8,51)$ 0.91 [0.5194]
M4	0.258%	–8.790	–8.911	–8.991	[0.459]	$F(4,55)$ 4.25^{**} [0.0046]	$F(8,51)$ 2.14^{*} [0.0481]
M5	0.245%	–8.848	–8.990	–9.082	[0.791]	$F(3,55)$ 2.90^{*} [0.0430]	$F(7,51)$ 1.33 [0.2557]
M6	0.256%	–8.906	–8.988	–9.040	[0.528]	$F(6,55)$ 2.97^{*} [0.0138]	$F(10,51)$ 1.80^{+} [0.0852]
M7	0.235%	–8.889	–9.051	**–9.156**	[0.968]	$F(2,55)$ 1.17 [0.3170]	$F(6,51)$ 0.55 [0.7685]
M8	0.243%	**–8.958**	–9.059	–9.125	[0.823]	$F(5,55)$ 1.98^{+} [0.0963]	$F(9,51)$ 1.16 [0.3403]
M9	0.246%	–8.842	–8.984	–9.076	[0.775]	$F(3,55)$ 3.02^{*} [0.0374]	$F(7,51)$ 1.38 [0.2346]
M10	0.251%	–8.661	–8.863	–8.995	— — —	— — —	$F(4,51)$ 2.25^{+} [0.0769]

Notes:

1. Models M1–M9 are defined in Table 6. Model M10 is equation (14) in de Brouwer and Ericsson (1998). Model M11 is the minimal union model of Models M1–M9. Model M12 is the minimal union model of Models M1–M9 and equation (14) in de Brouwer and Ericsson (1998).
2. The minimum values of the SC, HQ, and AIC are in bold.
3. F-statistics are reported in the format given in Table 2.

6 The Selected Model

This section summarizes the statistical and economic properties of the finally selected model, i.e., model M8. This model is:

$$\begin{aligned}\Delta p_t \quad = \quad & + \underset{(0.0054)}{0.0155}\ \Delta pet_t - \underset{(0.0098)}{0.0266}\ \Delta_4 i_t \\ & - \underset{(0.0047)}{0.0840}\ ecm_{t-1} + \underset{(0.0026)}{0.0110}\ D_t - \underset{(0.0434)}{0.7563}\end{aligned} \tag{10}$$

$T = 65$ [1977Q3–1993Q3] $\quad R^2 = 0.86 \quad \hat{\sigma} = 0.243\%$

$DW = 2.27 \quad AR : F(5, 55) = 1.44$ [0.2251]

$ARCH : F(4, 52) = 0.38$ [0.8216] $\quad Normality : \chi(2) = 1.16$ [0.5594]

$HeteroA : F(7, 52) = 0.57$ [0.7754] $\quad HeteroB : F(10, 49) = 0.68$ [0.7411]

$RESET : F(1, 59) = 0.05$ [0.8219].

Equation (10) reports a battery of diagnostic test statistics, and none of these statistics rejects at standard levels. The residuals appear to be white noise (AR), normally distributed (*Normality*), and homoscedastic ($ARCH$, *HeteroA*, *HeteroB*, and *RESET*).

Figure 5 plots the actual and fitted values for equation (10), the corresponding residuals, the histogram and estimated density of the residuals, and the residual autocorrelogram. Visually, the residuals exhibit no unusual properties. *Inter alia*, the residuals' estimated density is very close to a normal density, and the residual autocorrelations are all close to zero.

Tests of parameter constancy. Figures 6, 7, and 8 respectively report the recursive estimates, the recursive t-ratios, and the various recursive test statistics for equation (10). From Figure 6, all of the recursive estimates are both numerically and statistically constant over time, with the estimated coefficients typically fluctuating only a few tenths of their *ex ante* estimated standard errors. The recursive t-ratios in Figure 7 generally grow steadily over time or decline steadily over time, depending upon whether the coefficient is positive or negative. One notable exception is the t-ratio on Δpet_t: that t-ratio increases sharply at the time of the Gulf War. The associated change in petrol prices provides considerable additional information about the corresponding coefficient over a relatively short time. See Campos and Ericsson (1999) for further analysis of the accrual of information over time.

The Chow statistics in Figure 8 confirm the empirical constancy of equation (10) that is visually apparent from that equation's recursive estimates. Only a single one-step Chow statistic is statistically significant at the 5% level, and that statistic is only barely significant at that level. None of the breakpoint Chow statistics is significant at the 5% level. That is, no split of the sample obtains a rejection of constancy. None of the forecast Chow statistics is significant at the 5% level, either.

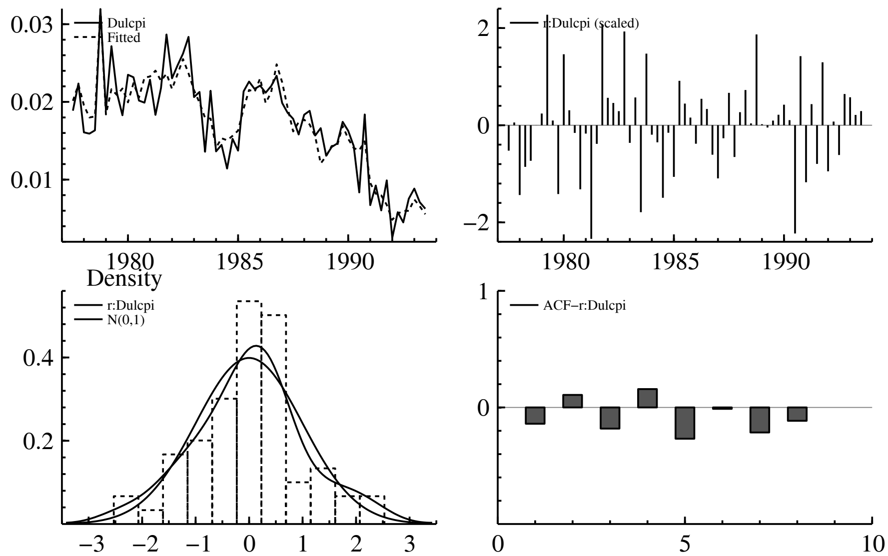

Figure 5: Actual and fitted values, residuals, the histogram and estimated density of the residuals, and the residual autocorrelogram for model M8.

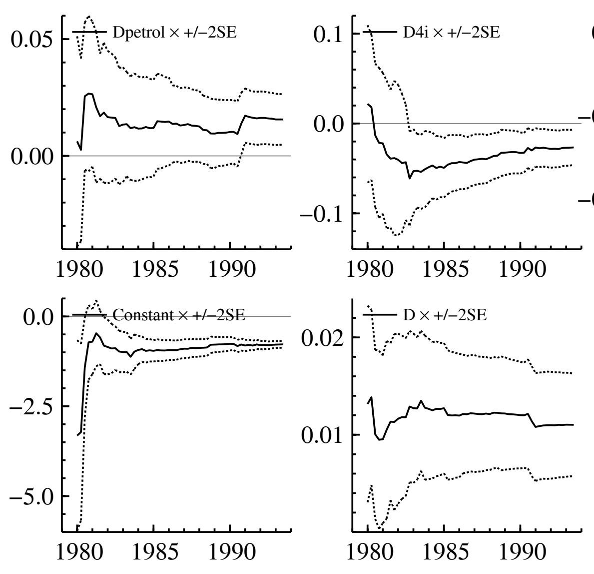

Figure 6: Recursive estimates for the coefficients in model M8.

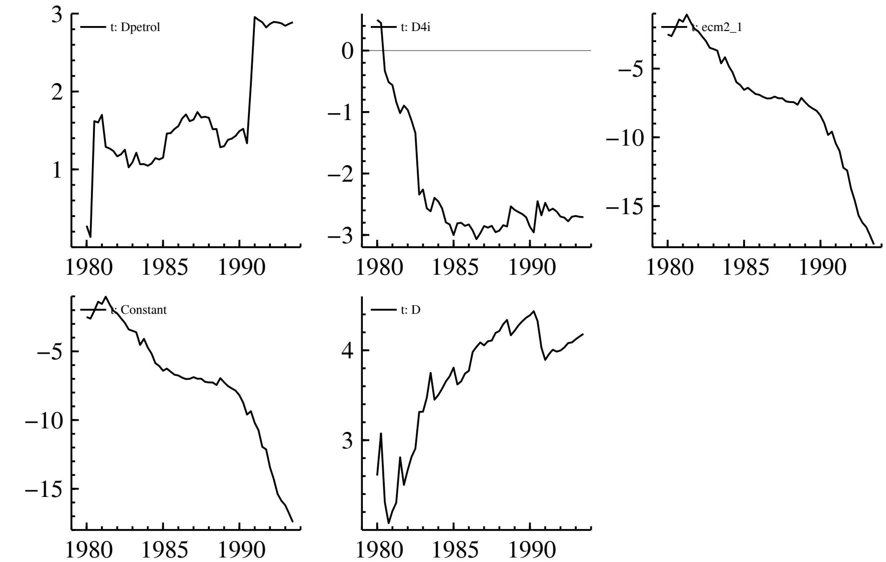

Figure 7: Recursive t-ratios for the coefficients in model M8.

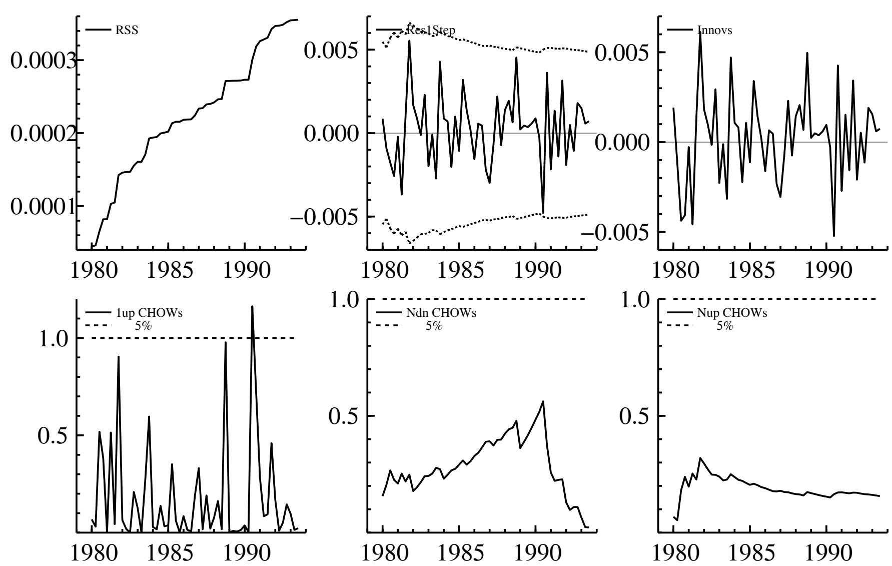

Figure 8: The recursive RSS, one-step residuals, standardized innovations, and one-step, breakpoint, and forecast Chow statistics for model M8.

The diagnostic statistics and the statistics for testing parameter constancy all point to equation (10) being a well-specified, empirically constant, congruent model. Furthermore, equation (10) improves upon the design of prior models on this dataset.

7 Conclusions

This paper advances de Brouwer and Ericsson's (1998) model of Australian inflation in three directions: cointegration analysis, treatment of weak exogeneity, and model design. Computer-automated model selection with the software package PcGets helps obtain a more parsimonious, empirically constant, data-coherent, encompassing error correction model for inflation in Australia. The level of consumer prices is a mark-up over domestic and import costs, with adjustments for dynamics and relative aggregate demand.

Several general remarks are germane, and each suggests possible extensions to the current analysis. First, the empirical ambiguity between some of the final models may be resolved by additional data: specifically, by more recent data. Dwyer and Leong (2001, Appendix D) report an Australian inflation equation similar to the one in de Brouwer and Ericsson (1998), but estimated over 1985Q1–2000Q1 rather than 1977Q3–1993Q3. A combined dataset may provide evidence that neither dataset contains individually. Second, improvements to the algorithm in PcGets may obtain an improved model specification. Computer-automated model selection algorithms are still in their youth—if not in their infancy—and considerable analytical and Monte Carlo research is ongoing; see Hendry and Krolzig (1999, 2003a, 2003b), Krolzig and Hendry (2001), and Hoover and Perez (2000). Third, other researchers may have insights on the empirical modeling of Australian inflation—insights not yet gleaned by the current researchers. The resulting model improvements would then continue a progressive research strategy pursued by the current paper.

References

Akaike, H. (1973) "Information Theory and an Extension of the Maximum Likelihood Principle", in B. N. Petrov and F. Csaki (eds.) *Second International Symposium on Information Theory*, Akademiai Kiado, Budapest, 267–281.

Akaike, H. (1981) "Likelihood of a Model and Information Criteria", *Journal of Econometrics*, 16, 1, 3–14.

Atkinson, A. C. (1981) "Likelihood Ratios, Posterior Odds and Information Criteria", *Journal of Econometrics*, 16, 1, 15–20.

Boughton, J. M. (1993) "The Demand for M1 in the United States: A Comment on Baba, Hendry and Starr", *Economic Journal*, 103, 420, 1154–1157.

de Brouwer, G., and N. R. Ericsson (1995) "Modelling Inflation in Australia", International Finance Discussion Paper No. 530, Board of Governors of the Federal Reserve System, Washington, D.C., November; and Research Discussion Paper No. 9510, Reserve Bank of Australia, Sydney, Australia, November.

de Brouwer, G., and N. R. Ericsson (1998) "Modeling Inflation in Australia", *Journal of Business and Economic Statistics*, 16, 4, 433–449.

Campos, J., and N. R. Ericsson (1999) "Constructive Data Mining: Modeling Consumers' Expenditure in Venezuela", *Econometrics Journal*, 2, 2, 226–240.

Doornik, J. A., and D. F. Hendry (2001a) *GiveWin Version 2: An Interface to Empirical Modelling*, Timberlake Consultants Press, London.

Doornik, J. A., and D. F. Hendry (2001b) *Modelling Dynamic Systems Using PcGive 10*, Volume 2, Timberlake Consultants Press, London.

Duesenberry, J. (1950) "The Mechanics of Inflation", *Review of Economics and Statistics*, 32, 144–149.

Durevall, D. (1998) "The Dynamics of Chronic Inflation in Brazil, 1968–1985", *Journal of Business and Economic Statistics*, 16, 4, 423–432.

Dwyer, J., and K. Leong (2001) "Changes in the Determinants of Inflation in Australia", Research Discussion Paper No. 2001–02, Reserve Bank of Australia, Sydney, Australia, May.

Ericsson, N. R., D. F. Hendry, and H.-A. Tran (1994) "Cointegration, Seasonality, Encompassing, and the Demand for Money in the United Kingdom", Chapter 7 in C. P. Hargreaves (ed.) *Nonstationary Time Series Analysis and Cointegration*, Oxford University Press, Oxford, 179–224.

Franz, W., and R. J. Gordon (1993) "German and American Wage and Price Dynamics: Differences and Common Themes", *European Economic Review*, 37, 4, 719–762 (with discussion).

Hannan, E. J., and B. G. Quinn (1979) "The Determination of the Order of An Autoregression", *Journal of the Royal Statistical Society, Series B*, 41, 190–195.

Hendry, D. F. (1995) "On the Interactions of Unit Roots and Exogeneity", *Econometric Reviews*, 14, 4, 383–419.

Hendry, D. F., and J. A. Doornik (1994) "Modelling Linear Dynamic Econometric Systems", *Scottish Journal of Political Economy*, 41, 1, 1–33.

Hendry, D. F., and J. A. Doornik (2001) *Empirical Econometric Modelling Using PcGive 10*, Volume 1, Timberlake Consultants Press, London.

Hendry, D. F., and H.-M. Krolzig (1999) "Improving on 'Data Mining Reconsidered' by K. D. Hoover and S. J. Perez", *Econometrics Journal*, 2, 2, 202–219.

Hendry, D. F., and H.-M. Krolzig (2001) *Automatic Econometric Model Selection Using PcGets 1.0*, Timberlake Consultants Press, London.

Hendry, D. F., and H.-M. Krolzig (2002) "New Developments in Automatic General-to-specific Modeling", in B. P. Stigum (ed.) *Econometrics and the Philosophy of Economics*, MIT Press, Cambridge, Massachusetts, forthcoming.

Hendry, D. F., and H.-M. Krolzig (2003) "The Properties of Automatic *Gets* Modeling", mimeo, Nuffield College, University of Oxford, Oxford, England, June.

Hendry, D. F., and R. M. Starr (1993) "The Demand for M1 in the USA: A Reply to James M. Boughton", *Economic Journal*, 103, 420, 1158–1169.

Hoover, K. D., and S. J. Perez (1999) "Data Mining Reconsidered: Encompassing and the General-to-specific Approach to Specification Search", *Econometrics Journal*, 2, 167–191 (With discussion).

Hoover, K. D., and S. J. Perez (2000) "Truth and Robustness in Cross-country Growth Regressions", mimeo, Economics Department, University of California at Davis, Davis, California.

Hylleberg, S. (ed.) (1992) *Modelling Seasonality*, Oxford University Press, Oxford.

Johansen, S. (1988) "Statistical Analysis of Cointegration Vectors", *Journal of Economic Dynamics and Control*, 12, 2/3, 231–254.

Johansen, S. (1991) "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models", *Econometrica*, 59, 6, 1551–1580.

Johansen, S. (1995) *Likelihood-based Inference in Cointegrated Vector Autoregressive Models*, Oxford University Press, Oxford.

Johansen, S., and K. Juselius (1990) "Maximum Likelihood Estimation and Inference on Cointegration—With Applications to the Demand for Money", *Oxford Bulletin of Economics and Statistics*, 52, 2, 169–210.

Juselius, K. (1992) "Domestic and Foreign Effects on Prices in an Open Economy: The Case of Denmark", *Journal of Policy Modeling*, 14, 4, 401–428.

Krolzig, H.-M., and D. F. Hendry (2001) "Computer Automation of General-to-specific Model Selection Procedures", *Journal of Economic Dynamics and Control*, 25, 6–7, 831–866.

Metin, K. (1998) "The Relationship Between Inflation and the Budget Deficit in Turkey", *Journal of Business and Economic Statistics*, 16, 4, 412–422.

Reserve Bank of Australia (1994a) *1994 Report and Financial Statements*, Reserve Bank of Australia, Sydney, Australia.

Reserve Bank of Australia (1994b) "Measuring 'Underlying' Inflation", *Reserve Bank of Australia Bulletin*, August, 1–6.

Richards, T., and G. Stevens (1987) "Estimating the Inflationary Effects of Depreciation", Research Discussion Paper No. 8713, Reserve Bank of Australia, Sydney, Australia.

Schwarz, G. (1978) "Estimating the Dimension of a Model", *Annals of Statistics*, 6, 461–464.