

The Methodology and Practice of Econometrics

A Festschrift in Honour of
David F. Hendry

Edited by
Jennifer L. Castle
and
Neil Shephard

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© Jennifer L. Castle & Neil Shephard 2009

The moral rights of the authors have been asserted
Database right Oxford University Press (maker)

First published 2009

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloging-in-Publication Data

The methodology and practice of econometrics : a festschrift in honour of
David F. Hendry / Jennifer Castle and Neil Shephard.

p. cm.

Includes bibliographical references.

ISBN 978-0-19-923719-7 (acid-free paper)

1. Econometrics. 2. Econometric models. I. Hendry, David F. II. Castle,
Jennifer, 1979- III. Shephard, Neil.

HB139.M474 2009

330.01'5195—dc22

2008055210

Typeset by SPI Publisher Services, Pondicherry, India

Printed in Great Britain

on acid-free paper by

CPI Antony Rowe, Chippenham, Wiltshire

ISBN 978-0-19-923719-7

1 3 5 7 9 10 8 6 4 2

2

Empirical Identification of the Vector Autoregression: The Causes and Effects of US M2

Kevin D. Hoover, Selva Demiralp, and Stephen J. Perez*

A commonplace in economics is that theory and empirics stand in a position of mutual support.¹ But there can be little doubt that for some time theory has had the upper hand. We see this in the widespread view that it is only through a priori theory that a system of equations can be identified, and we see it in the equally widespread view that data-mining—that is, using empirical data to shape the formulation of an econometric model—is, in all of its forms, an unequivocally bad thing. We are left with the odd view that theory proposes and data disposes. On this principle, empirical evidence is, at best, destructive, but cannot be constructive. In reality we find it hard to live up to such fine principles, but the economics profession naturally tries its best to hide (or disguise) its apostasy.

David Hendry belongs to the minority camp that embraces the notion that empirical investigation plays a *constructive* role in developing our understanding of the economy. In his inaugural address at the London School of Economics, 'Econometrics: Alchemy or Science?' (1980), Hendry cogently made the case for constructive econometrics—a chemistry, not an alchemy. Hendry and the so-called LSE school (Gilbert, 1986; Mizon, 1995) recognize that, for many important economic problems, economic theory is too weak—and absent learning more about the facts on the ground—always will be too

* The authors thank Søren Johansen for valuable econometric advice.

¹ Kant is frequently quoted in support without a specific reference as: 'Theory without empirics is empty. Empirics without theory is blind.' or 'Experience without theory is blind, but theory without experience is mere intellectual play.' Both seem to be overly free translations of '*Gedanken ohne Inhalt sind leer, Anschauungen ohne Begriffe sind blind*' ('Thoughts without contents are empty; opinions without concepts are blind'—Hoover translation); Kant (1787, part 2, section 1).

weak to fulfill the dominant role that the reigning ideology of economics assigns it. If one wants to know the truth about the economy, one often has to grub through the data with economic theory providing—at most—broad guidance and not detailed specifications. In his vital work on encompassing and general-to-specific specification search, culminating in his work on automatic model selection (particularly, in the PcGets software—see Hendry and Krolzig 1999, 2001; Krolzig and Hendry, 2001), Hendry grasps the nettle. Rather than hiding constructive econometrics to suit the prejudices of the profession, he asks sensibly, how should it be done well.

We have long taken Hendry's side of the debate. And indeed, we have contributed something to automatic model selection in the spirit of the LSE school (Hoover and Perez 1999, 2004). We have also long been interested in what, at first, might seem an orthogonal concern: causality in economics (Perez, 1998; Hoover, 2001; Demiralp and Hoover, 2003).² We believe that there is an important complementarity, which has yet to be fully exploited, between Hendry's general-to-specific approach and recent work on the graph-theoretic approach to causal modelling (Spirtes, Glymour, and Scheines, 2001; Pearl, 2000; Swanson and Granger, 1997; Demiralp and Hoover, 2003; Hoover, 2005). Krolzig (2003) demonstrates that PcGets is effective at recovering the dynamic structure of a system of equations (a structural vector autoregression or SVAR), provided that one starts with a diagonal covariance matrix—in other words, provided that one knows the contemporaneous causal order of the SVAR. Graph-theoretic causal-search algorithms can aid in the discovery of that causal order, so that, together with a general-to-specific search algorithm, we have some hope of identifying the structure of the SVAR empirically. (Instead of PcGets, we in fact use Autometrics—see the chapter by Doornik in this volume for description of the algorithm.)

In this chapter we provide a concrete illustration of the complementary use of graph-theoretic causal modelling and automated general-to-specific specification search. Our problem is to identify the factors determining the US M2 monetary aggregate and its role in the transmission of monetary policy—a problem for which economic theory provides only the broadest guidance.

2.1 Understanding M2

M2 consists of liquid deposits, small time deposits, retail money funds, and currency in circulation. Even though, owing to the widespread use of alternative financial market instruments, the relationship between monetary

² And on which Hendry might be seen as somewhat sceptical—see Hendry *et al.* (1990), p. 184.

aggregates and income growth has loosened over the last decade, the Federal Reserve still regards the pattern of M2 growth as providing information about the conditions of aggregate demand. M2 growth is monitored by the Monetary and Reserve Analysis Section of the Division of Monetary Affairs of the Board of Governors of the Federal Reserve System.³ Although the Section implicitly assumes a relatively rich causal structure in explaining the process of M2 growth, there are no formal studies (inside or outside the Federal Reserve) that analyse M2 growth analytically. The Federal Reserve's econometric models do forecast M2 growth, but these models are mostly driven by the quantity theory of money, and omit many of the implicit structural considerations that the Section regards as important.

Before each meeting of the Federal Open Market Committee (FOMC), the Section prepares a contribution to the briefing document, known as the 'Blue-book', in which it analyses the growth of M2 in relation to a number of factors that have not yet been investigated structurally. The economic theory used is fairly broad brush. According to the quantity theory of money, the growth rate of money should equal the growth rate of nominal income, adjusting for the trend in velocity. The Section's analysis, therefore, starts by anchoring the underlying growth rate of M2 to the growth rate of GDP, and then considers 'special factors' that may cause deviations of M2 growth rate from that of GDP. These special factors comprise:

- i) Interest rate effects: changes in the Federal funds rate target lead to subsequent changes in the opportunity cost of holding M2 type of assets.
- ii) Equity market effects: high volatility and downwards revisions to the expectations of earnings on equities earnings *ceteris paribus* boost M2 as investors substitute away from the stock market and into safe and liquid M2-type assets.
- iii) Other special factors including: activity in *mortgage-backed securities*, as mortgage servicers temporarily accumulate the proceeds of pre-paid mortgages in the liquid-deposits component of M2; *tax effects*, which influence the money-market-deposit-account (MMDA) component of liquid deposits as people pay their taxes out of their savings accounts; and *currency shipments abroad*.

To illustrate, Table 2.1 shows the growth rates of the components of M2 as of December 2004. Comparing the second and the seventh columns shows to that in the second quarter of 2004 M2 growth exceeded that of GDP growth (row 1), whereas in the third quarter it fell behind (row 2). The Section attributed the accelerated pace in the second quarter to: (i) mortgage

³ This chapter was conceived while Demiralp was employed as a staff economist in the Monetary and Reserve Analysis Section.

Table 2.1. Growth rates of M2 and its major components (percent per annum)

	M2	Liquid Deposits	Small Time Deposits	M2 Money Funds	Currency	Nominal GDP	Inflation ¹
2004:2	9.7	17.1	-4.6	-7.2	4.0	6.6	4.7
2004:3	2.7	4.5	1.8	-11.3	9.3	5.1	6.2
October 2004	2.5	6.1	3.3	-20.2	4.2		
November 2004	6.3	8.0	5.1	-5.7	9.9		

Source: Board of Governors of the Federal Reserve System and Bureau of Economic Analysis.

¹ Measured by the personal consumption expenditure deflator.

refinancing activity in April (which was boosted by a decline in mortgage interest rates), and (ii) inflows from equity and bond funds as well as increased deposits of tax refunds in May. Meanwhile the slowdown in the third quarter was mostly explained by the rising opportunity cost in the face of a series of steps to tighten monetary policy. In each case, the largest component of M2, liquid deposits, accounts for most of the overall growth rate (see Table 2.2).

Liquid Deposits, which constitute about 65% of M2, comprise *demand deposit accounts* (DDAs), *other certificates of deposits* (OCDs), and *savings deposits* (including MMDAs). DDAs and OCDs are the most liquid of the components of liquid deposits, and appear to respond to changes in the opportunity cost of holding M2 and similar assets. In addition to this opportunity-cost channel, running from changes in the Federal funds rate to the opportunity cost of M2 to liquid deposits, a decline in the Federal-funds-rate target may also lead to a decline in mortgage rates, a consequent rise in the mortgage refinancing, and a rise in liquid deposits to meet the temporary need of mortgage servicers to park funds for several weeks until the mortgage-backed securities are redeemed. Other transitory changes in the holdings of liquid deposits may be related to tax payments, influencing especially DDAs and OCDs. The MMDA component of liquid deposits is a close substitute for stock

Table 2.2. Contributions of M2 components to M2 growth

	2004:2	2003:3	October 2004	November 2004
M2	9.7	2.7	2.5	6.3
Liquid Deposits	10.8	2.9	4.0	5.2
MMMF	-0.9	-1.4	-2.3	-0.6
Small Time	-0.6	0.2	0.4	0.6
Currency	0.4	1.0	0.5	1.1

Source: Board of Governors of the Federal Reserve System.

Components may not sum to total because traveller's cheques are not shown.

mutual funds and may, therefore, display sensitivity to the performance of the stock market. Events, such as domestic or international political crises, also boost the demand for safe and liquid M2 components. On the other hand, steepening of the yield curve, because of an increase in long-term yields or a looser monetary policy, may reduce the growth of liquid deposits as investors substitute into longer-term assets.

After a brief detour to set out the strategy of empirical investigation, we will in sections 2.3 through 2.6 investigate to what degree the data support qualitatively and quantitatively the Monetary and Reserve Analysis Section's informal understanding of the role of M2 in the transmission of monetary policy.

2.2 Empirical Identification

Our approach will be to specify a structural vector autoregression as far as possible using the tools of graph-theoretic causal-search algorithms and Doornik and Hendry's Autometrics software. Autometrics is sufficiently well described (Doornik, this volume) that we need not spend time on explaining the principles of its operation. Although the same is not true of the graph-theoretic causal-search algorithms, we nonetheless will give only an informal sketch and refer the reader to the fuller descriptions cited in the introduction.

The SVAR can be written as:

$$\mathbf{A}_0 \mathbf{Y}_t = \mathbf{A}(L) \mathbf{Y}_{t-1} + \mathbf{E}_t, \quad (2.1)$$

where \mathbf{Y}_t is an $n \times 1$ vector of contemporaneous variables, \mathbf{A}_0 is a full rank $n \times n$ matrix with ones on the main diagonal and possibly non-zero off-diagonal elements; $\mathbf{A}(L)$ is a polynomial in the lag operator, L ; and \mathbf{E}_t is an identical independent normal $n \times 1$ vector of error terms $\mathbf{E}_t \sim N(0, \Sigma)$. Let $\mathbf{E} = [\mathbf{E}_t]$, $t = 1, 2, \dots, T$, then the covariance matrix $\Sigma = E(\mathbf{E}\mathbf{E}')$ is diagonal. The individual error terms (shocks) can be assigned unequivocally to particular equations because Σ is diagonal. The matrix \mathbf{A}_0 defines the causal interrelationships among the contemporaneous variables. The system is identified provided that there are $n(n-1)/2$ zero restrictions on \mathbf{A}_0 . For any just-identified system, \mathbf{A}_0 can be rendered lower triangular by selecting the appropriate order of the variables \mathbf{Y} along with the conformable order the rows of \mathbf{A}_0 . This is the *recursive* (or *Wold causal*) order.

Starting with the SVAR as the data-generating process (DGP), premultiplying by \mathbf{A}_0^{-1} yields the reduced-form or ordinary vector autoregression (VAR):⁴

$$\mathbf{Y}_t = \mathbf{A}_0^{-1} \mathbf{A}(L) \mathbf{Y}_{t-1} + \mathbf{A}_0^{-1} \mathbf{E}_t = \mathbf{B}(L) \mathbf{Y}_{t-1} + \mathbf{U}_t. \quad (2.2)$$

⁴ This and the next six paragraphs are closely based on Hoover (2005).

The transformed error terms are $\mathbf{U} = [\mathbf{U}_t]$, $t = 1, 2, \dots, T$. And, in general, $\mathbf{\Omega} = E(\mathbf{U}\mathbf{U}')$ is, unlike $\mathbf{\Sigma}$, not diagonal. If we knew \mathbf{A}_0 , then recovery of the SVAR (equation 2.1) from the easily estimated VAR (equation 2.2) would be straightforward. There are, however, a large number of $n \times n$ matrices, \mathbf{P}_i that may be used to premultiply equation (2.2) such that the covariance matrix $\mathbf{\Omega} = E(\mathbf{P}_i\mathbf{U}(\mathbf{P}_i\mathbf{U}'))$ is diagonal. The central identification problem for SVARs is to choose the one member of \mathbf{P}_i that corresponds to the data-generating process, that is to find $\mathbf{P}_i = \mathbf{A}_0$, when \mathbf{A}_0 is unknown.

Identification requires at least $n(n-1)/2$ restrictions on \mathbf{P}_i . If we restrict ourselves to zero restrictions on recursive systems, then any just-identified \mathbf{P}_i can be arranged in the form of one of the $n!$ Choleski orderings (or decompositions) corresponding to each of the possible permutations of the variables in \mathbf{Y} . For any given permutation, there is a unique lower triangular \mathbf{P}_i such that $\mathbf{P}_i\mathbf{\Omega}\mathbf{P}_i'$ is diagonal. The just-identified SVARs specified by each of the Choleski orderings is observationally equivalent in the sense that they each have the same likelihood. Yet, the different orderings have different consequences for the dynamics of the SVAR.

On what basis, then, should we choose? If we are restricted to just-identified SVARs, then we have little choice but to appeal to economic theory to tell us what the causal order should be. This is, in fact, what almost all practitioners of VAR methodologies profess to do. Unfortunately, formal economic theory is rarely decisive about causal order. So, VAR practitioners either choose the order arbitrarily, sometimes with an accompanying claim that their results are robust to alternative causal orderings—apparently unaware that such robustness arises in just those cases that the contemporaneous terms are unimportant. Or they appeal, not so much to theory, as to ‘just so’ stories: intuition or commonsense tells them that, say, financial markets adjust more quickly than goods markets, so that interest rates, for instance, ought to be causally ordered ahead of real GDP. It is usually easy, however, to tell a ‘just so’ story to justify almost any order—the time order of variables that are contemporaneously related at the given frequency of observation being especially unreliable. There is a special irony that this strategy should be so commonly accepted among VAR practitioners. After all, Sims’s (1980) motivation in initiating the VAR program was to avoid the need to appeal to ‘incredible’ identifying restrictions.

If, however, the true SVAR is overidentified, then we have another option. Graph-theoretic causal search provides a method of choosing \mathbf{P}_i , very much in the spirit of Hendry’s general-to-specific model selection. In a causal graph, arrows connecting causal variables to their effects represent causal relationships. Spirtes *et al.* (2001) and Pearl (2000) show that there are isomorphisms between graphs and the probability distributions of variables. In particular, certain graphical patterns imply certain relationships of conditional independence and dependence among the variables. The graph of the DGP can also

be represented through the restrictions on \mathbf{A}_0 . Working backwards from statistical measures of conditional independence and dependence, it is possible to infer the class of graphs compatible with the data. Sometimes that class has only a single member, and then \mathbf{A}_0 can be identified statistically.

The key ideas of the graph-theoretic approach are simple. Suppose that $A \rightarrow B \rightarrow C$ (that is, A causes B causes C). A and C would be dependent, but conditional on B , they would be independent. Similarly for $A \leftarrow B \leftarrow C$. In each case, B is said to *screen* A from C . Suppose that $A \leftarrow B \rightarrow C$. Then, once again A and C would be dependent, but conditional on B , they would be independent. B is said to be the *common cause* of A and C . Now suppose that A and B are independent conditional on sets of variables that exclude C or its descendants, and $A \rightarrow C \leftarrow B$, and none of the variables that cause A or B directly causes C . Then, conditional on C , A and B are dependent. C is called an *unshielded collider* on the path ACB . (A *shielded collider* would have a direct link between A and B .)

Causal search algorithms use a statistical measure of independence, commonly a measure of conditional correlation, to check systematically the patterns of conditional independence and dependence and to work backwards to the class of admissible causal structures. In this chapter, we use the PC algorithm, the most common of the causal-search algorithms (Spirtes *et al.*, 2001, pp. 84–85; Pearl, 2000, pp. 49–51; Cooper, 1999, p. 45, Figure 22). It assumes that graphs are *acyclical* or strictly recursive—that is, loops in which $A \rightarrow B \rightarrow C \rightarrow A$ are ruled out. Naturally, acyclicity also rules out simultaneity—that is, a very tight loop in which $A \rightarrow B \rightarrow A$ (or $A \leftrightarrow B$). While the assumption of acyclicity is restrictive, it is nonetheless more general than limiting SVARs to Choleski orders, which remain the default in most VAR studies.

The details of the PC algorithm are described in Demiralp and Hoover (2003). Essentially, it begins with the complete set of variables in the VAR densely connected by undirected edges, represented as lines in a graph without arrowheads. It then tests for unconditional correlations and removes any uncorrelated edges. It then tests for correlations conditional on one other variable, again removing edges for which correlations vanish on conditioning. It then proceeds to conditioning on two, three, ... variables. The result is an undirected *skeleton*. The algorithm begins orienting edges by seeking triples of linked variables ($A-B-C$) in which the variables on the endpoints (A and C) are independent on some conditioning set, but become dependent when conditioning on the intermediate variable (B). This is the pattern of an unshielded collider, and the edges are then oriented ($A \rightarrow B \leftarrow C$). Some edges may be oriented *logically* (rather than statistically), based on maintaining the assumption of acyclicity and avoiding implying the existence of unshielded colliders not identified statistically.

Not all causal graphs are recoverable from the probability distribution. Graphs that have the same unshielded colliders and the same skeleton are

observationally equivalent (Pearl, 2000, p. 19). If the true graph is a member of an observationally equivalent set, the algorithm will not orient the edges that distinguish one member of the set from another. In these cases, unoriented edges can be oriented in either direction without changing the likelihood, provided that no new unshielded colliders or cyclicalities are introduced. Also, the maintained assumption of acyclicity notwithstanding, the algorithm will sometimes identify edges as bidirectional as a result of either ambiguity in the statistical test because of small samples or omitted latent variables.

Following Swanson and Granger (1997), we treat the estimated errors (\hat{U}_t) from the VAR in equation (2.2) as the original data purged of their dynamics. The covariance matrix of these transformed data ($\hat{\Omega}$) provides the necessary data for computing the various conditional correlations required by the PC algorithm. The algorithm selects a graph that best represents the causal order, and this graph in turn corresponds to particular zeros in (and overidentifying restrictions on) \mathbf{A}_0 .

Demiralp and Hoover (2003) provide Monte Carlo evidence that shows that the PC algorithm is highly effective at recovering the skeleton of the DGP graph and moderately effective at recovering the directions of individual links, provided that signal-to-noise ratios are high enough. Demiralp, Hoover, and Perez (2008) develop and validate a bootstrap procedure to assess the effectiveness of the closely related SGS algorithm. The procedure constructs many simulations of the VAR, equation (2.2), based on the actual coefficient estimates ($\hat{\mathbf{B}}(L)$) and resampling of the columns of \hat{U}_t , runs the search algorithm, and keeps track of the distribution of edges in the resulting graphs. The bootstrap method is essentially heuristic and provides guidance for more formal investigation of the overidentifying restrictions on \mathbf{A}_0 .

Once we have selected \mathbf{A}_0 as the orthogonalizing transformation to transform the VAR, equation (2.2), into the SVAR, equation (2.1), then we can appeal to Krolzig's (2003) evidence for the effectiveness of PcGets at locating the true restrictions on the lagged coefficients—that is, the placement of zeros in the matrix $\mathbf{A}(L)$.

2.3 Data

The data consist of eleven monthly series that run from 1990:02 to 2005:03. Sources and details are provided in Appendix A.⁵ Our main interest is in M2 and its role in the transmission mechanism. M2 is represented by its active component, (the logarithm of) liquid deposits (*LIQDEP*). Following

⁵ The data are available from <http://www.econ.duke.edu/~kdh9/research.html>.

the considerations of the Monetary and Reserve Analysis Section discussed in section 2.1, the principal factors related to liquid deposits are core CPI inflation (*COREINF*) and a monthly proxy for real GDP, (the logarithm of) industrial production (*IP*). The additional considerations of the equity market are represented by (the logarithm of) the S&P 500 stock market index (*SP500*), its price-earnings ratio (*SPPE*), and stock market volatility (*VOL*). Mortgage activity is represented by (the logarithm of) an index of mortgage refinancing (*REFI*) and the interest rate on 30-year fixed-rate mortgages (*MORG30*). Monetary policy is represented by the Federal funds rate (*FF*). The Section also monitors the opportunity cost of M2 (*M2OC*), which is constructed from the 3-month Treasury bill rate (*TBILL3*) and the own-rate on M2 (*M2OWN*). One question to be addressed is whether the two interest rates from which the opportunity cost of M2 is constructed enter only through *M2OC* or in fact have differential effects on other variables. The interest rates in the data set are rich enough to allow us to assess the interest-rate channel and the role of the yield curve in the transmission mechanism for monetary policy.⁶

As a preliminary, the data were graphed and tested for nonstationarity. Dickey-Fuller tests (with a constant and a trend), indicate that each of the series is very likely I(1), although for industrial production the test was borderline, and the series may even be better described as I(2).

2.4 Contemporaneous Causal Order

In order to test whether the opportunity cost of M2 is a satisfactory summary of the effects of its component rates, we test the restriction that the 10-variable VAR using *M2OC* is a valid restriction of the 11-variable VAR using *M2OWN* and *TBILL3*. The test is a likelihood ratio test with 40 restrictions: $\chi^2(40) = 476$, where the critical value $\chi^2_{0.05} = 56$. The 10-variable VAR is decisively rejected, and we analyse the 11-variable VAR hereafter.

Our first task in establishing the contemporaneous causal order among our variables is to estimate the unrestricted VAR—that is, equation (2.2), where $\mathbf{Y}_t = [\text{COREINF}, \text{FF}, \text{IP}, \text{SPPE}, \text{M2OWN}, \text{LIQDEP}, \text{SP500}, \text{REFI}, \text{MORG30}, \text{VOL}, \text{TBILL3}]'$. The corrected Akaike Information Criterion of Hurwicz and Tsai (1991), which seems well-adapted to this problem, selects a lag length of one. Both the Hannan-Quinn and Schwarz criteria also select only a single lag. However, all of these tests impose the same lag length on all equations. And, since in a later step, we intend to search for parsimonious, variable-specific dynamics, we should not be too restrictive at this stage. Therefore, we set

⁶ We have restricted ourselves to factors considered by the Section. There are other possible channels of monetary transmission. These are to some extent captured by *LIQDEP*, which is linked to the balance sheets of banks. A richer specification of the SVAR is possible, but it is beyond the scope of the present chapter.

the lag length to four, which will allow three lagged differences when we later construct error-correction specifications. We estimate the VAR and obtain the covariance matrix of \hat{U}_t , using it as the input into the PC algorithm, with a critical value of 10% for tests of conditional correlation. This is the critical value suggested on the basis of Monte Carlo studies by Spirtes *et al.* (1994, pp. 103–107) for the number of available observations (178 after accounting for lags). A critical value higher than the more common 5% is also justified by our concern not to restrict the specification too much (that is, we choose a ‘liberal’ strategy that shifts the balance somewhat towards the avoidance of type 2 error).

The algorithm selects the graph in Figure 2.1. While one should not read too much into the contemporaneous structure, since many important causal channels may operate with a lag, the graph is striking in that the variables are distinctly grouped: 1) the various interest rates are ordered as block recursively above the financial-asset variables; 2) core inflation is isolated contemporaneously from all other variables; and 3) mortgage refinancing is causally connected to the 30-year mortgage rate, but the direction of causation is

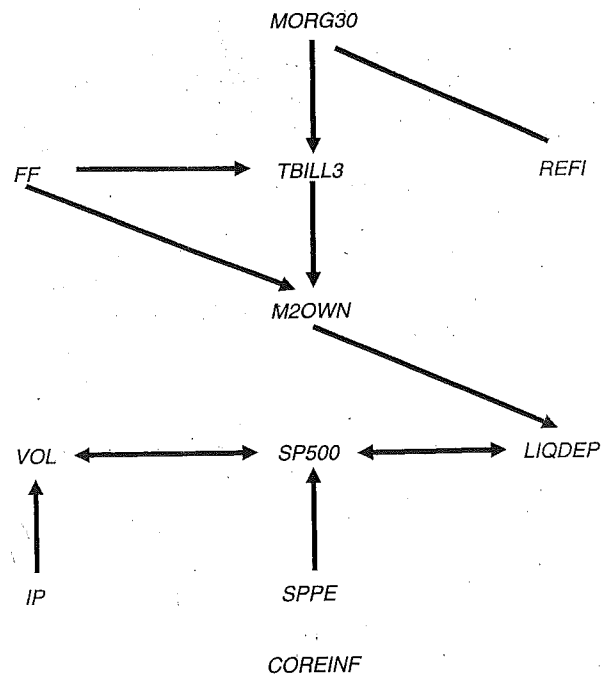


FIG. 2.1. Initial contemporaneous causal graph.

unresolved. Note the bidirectional edges between *VOL* and *SP500* and between *SP500* and *LIQDEP*.

How reliable is the identified graph? To evaluate it, we apply the bootstrap procedure of Demiralp, Hoover, and Perez (2008) with 10,000 replications. The results are shown in Table 2.3. There are 55 possible edges among eleven variables. The first three columns show the results of the PC algorithm for 16 of them—the 10 selected by the algorithm and 6 more that are selected in 10% or more of the bootstrap replications. The next five columns show the actual distribution of edges from the bootstrap. The last three columns present summary statistics: *exists* is the fraction of replications in which some edge is found (= 100 – *no edge*); *directed* is the percentage of edges discovered that have a definite direction; *net direction* is the difference between the percentage of edges directed rightward (→) and edges directed leftward (←).

The bootstrap presents strong evidence in favour of the existence of the first six edges in Table 2.3 (*exists* greater than 95%). In four of the six cases, the bootstrap favours the direction selected by the PC algorithm. One edge is undirected (*REFI* — *MORG30*). The bootstrap directs it only 35% of the time

Table 2.3. Bootstrap evaluation of the causal graph

	Causal Order Selected by the PC Algorithm ¹	Edge Identification (percent of bootstrap realizations) ²					Summary Statistics for Bootstrap Distribution ³		
		—	←	no edge	→	↔	exists	directed	net direction
<i>REFI</i>	— <i>MORG30</i>	65	22	0	11	3	100	35	-11
<i>SP500</i>	← <i>SPPE</i>	17	52	0	11	20	100	83	-41
<i>FF</i>	→ <i>M2OWN</i>	47	14	0	38	1	100	53	25
<i>FF</i>	→ <i>TBILL3</i>	7	5	1	79	8	99	93	73
<i>VOL</i>	↔ <i>SP500</i>	8	20	5	28	39	95	92	8
<i>MORG30</i>	→ <i>TBILL3</i>	7	3	5	75	10	95	93	72
<i>LIQDEP</i>	↔ <i>SP500</i>	0	0	45	30	24	55	100	30
<i>IP</i>	→ <i>VOL</i>	2	1	47	37	14	53	97	37
<i>M2OWN</i>	← <i>TBILL3</i>	0	10	64	16	10	36	99	6
<i>LIQDEP</i>	← <i>M2OWN</i>	2	8	72	7	12	28	94	-1
<i>REFI</i>	no edge <i>TBILL3</i>	0	0	82	15	3	18	98	14
<i>REFI</i>	no edge <i>M2OWN</i>	3	6	83	3	4	17	80	-3
<i>IP</i>	no edge <i>SPPE</i>	1	0	85	7	7	15	94	7
<i>SP500</i>	no edge <i>M2OWN</i>	1	4	86	1	9	14	95	-3
<i>COREINF</i>	no edge <i>SP500</i>	1	0	87	10	2	13	95	10
<i>IP</i>	no edge <i>M2OWN</i>	0	0	88	5	7	12	99	5

¹ 16 of 55 candidate edges; only edges that are identified as existing in 12% or more of the bootstrap replications are shown.

² Values indicate percentage of 10,000 bootstrap replications in which each type of edge is found. Based on the procedure in Demiralp, Hoover, and Perez (2007) with critical value of 2.5% for tests of conditional correlation (corresponding to the 10% critical value used in the PC algorithm).

³ *exists* = the percentage of bootstrap replications in which an edge is selected (= 100 – *no edge*); *directed* = edges directed as a percentage of edges selected; *net direction* = difference between edges directed right (→) and left (←).

with an 11-point preference for the order $REFI \leftarrow MORG30$. Another one of the six, $VOL \leftrightarrow SP500$ is bidirected. The bootstrap finds bidirectional edge in a plurality of cases, while in the others directs it $VOL \rightarrow SP500$ with an 8-point preference. The remainder of the edges chosen by the PC algorithm are less well supported by the bootstrap. The weakest case is for the edge $LIQDEP \leftarrow M2OWN$ with *exists* only registering at 28%. However, in all but one case, *net direction* agrees with the selected order. The exception is $M2OWN \leftarrow TBILL3$ for which the bootstrap would direct the edge in the opposite direction 6 points more frequently. We explore the problematic edges identified by the bootstrap more fully below.

A causal graph corresponds to a set of over-identifying restrictions (zero restrictions on A_0), which can be tested. The graph in Figure 2.1, however, cannot be tested as is, since it contains an undirected edge. There is, then, a two-member equivalence class one in which $REFI \rightarrow MORG30$, the other in which $REFI \leftarrow MORG30$. The overidentifying assumptions implied by the graph can be tested with either ordering of this edge; and, because they define an equivalence class, the result will be the same. The likelihood-ratio test of the overidentifying restrictions for the graph strongly rejects the restrictions (p -value of 0.002).

Since the risk, we believe, is greater of too tightly restricting the causal order, we investigate the graph further through an informal general-to-specific procedure. In Table 2.4, we investigate two search paths, one for each of the graphs in the equivalence class. Search I, with $REFI \rightarrow MORG30$, adds to the graph all of the edges that the bootstrap finds in 12% or more replications, directed as indicated by *net direction* in Table 2.3. As shown in Table 2.4, this model (General Model I) cannot be rejected against a just-identified SVAR (p -value = 0.14). The table also reports a sequence of tests in which successive edges are removed, starting with the edge with the lowest t -statistic when the SVAR based on the ordering starting with the general model is estimated. Each time an edge is removed, the p -value of the likelihood ratio test against the just-identified model is calculated. Test 1 results in a failure of the estimates to converge, so the edge ($LIQDEP \leftarrow SP500$) is restored and the edge with the next lowest t -statistic is removed. ($LIQDEP \leftarrow SP500$ is again removed with no converge problem in test 3.) In the end, six edges are removed, generating a sequence of specifications that cannot be rejected at a 10% critical value against the just-identified SVAR. Only the test 7 (omit $REFI \rightarrow TBILL3$) rejects. The remaining edges all correspond to statistically significant t -statistics. And the likelihood ratio test for the final graph against the just-identified SVAR cannot be rejected ($p = 0.102$).

Search II takes the graph selected by the PC algorithm with the undirected edge oriented as $REFI \leftarrow MORG30$. When supplemented with additional edges, using the same criterion as in Search I, it is easily rejected against the just-identified SVAR. Adding even more edges—all those found to exist in 5% or

Table 2.4. Contemporaneous causal structure: Specification search

	Specification	Likelihood Ratio Test against the Just-Identified Model (p -value)
Search I		
Initial Model	Graph in Figure 1 modified with $REFI \rightarrow MORG30$	0.002
General Model I	As above, plus: $REFI \rightarrow TBILL3$ $REFI \leftarrow M2OWN$ $IP \rightarrow SPPE$ $M2OWN \rightarrow SP500$ $IP \rightarrow M2OWN$ $COREINF \rightarrow REFI$	0.140
Tests of Restrictions		
1	omit $LIQDEP \leftarrow SP500$	no convergence
2	restore $LIQDEP \leftarrow SP500$; omit $VOL \rightarrow SP500$	0.163
3	omit $LIQDEP \leftarrow SP500$	0.185
4	omit $IP \rightarrow SPPE$	0.166
5	omit $M2OWN \leftarrow TBILL3$	0.146
6	omit $COREINF \rightarrow SP500$	0.123
7	omit $IP \rightarrow M2OWN$	0.102
8	omit $REFI \rightarrow TBILL3$	0.068
Search II		
Initial Model	Graph in Figure 1 modified with $REFI \leftarrow MORG30$	0.002
General Model II	As above, plus: all edges that appeared in more than 2.5% of bootstrap replications	0.069

even 2.5% of the bootstrap replications—still results in strong rejection of the specification. As a result, we discontinued further specification search for this graph and accepted the graph selected in Search I as our final graph.

Figure 2.2 shows the final graph. Notice that, compared to the initial graph (Figure 2.1), as well as orienting the undirected edge, it adds two edges ($REFI \leftarrow M2OWN$ and $M2OWN \rightarrow SP500$); it removes one edge ($M2OWN \leftarrow TBILL3$); and it turns the two bidirectional edges ($VOL \leftrightarrow SP500$ and between $SP500 \leftrightarrow LIQDEP$) into unidirectional edges.

2.5 The Lag Structure

Most VAR analysis would content itself with having established the contemporaneous causal order, which is all that is needed to identify independent shocks to the various equations in the SVAR, and then proceed to compute impulse-response functions and variance decompositions. We believe, however, that more is to be learned about the dynamic causal structure about

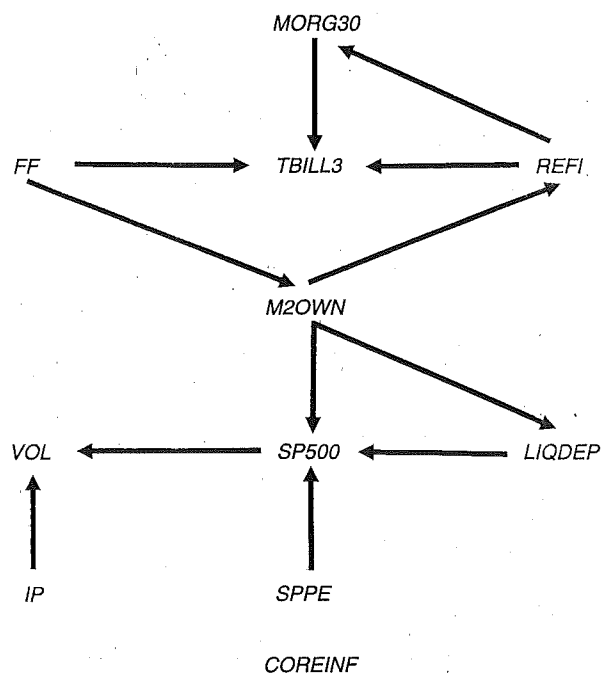


FIG. 2.2. Final contemporaneous causal graph.

which factors are truly important. And we believe that a more careful specification of the dynamics will deliver more precise estimates of the standard errors of impulse-response functions.

Our strategy is to follow the strategy outlined in Krolzig (2003) in which a general-to-specific search algorithm is used to select the lag structure of the model conditional on the contemporaneous causal structure. The initial selection of the contemporaneous causal structure is important, since equation-by-equation automated searches must start with a valid general unrestricted model in which the independent variables are either not the effects of the dependent variable or in which they are properly instrumented. The graph in Figure 2.2 gives just the needed information and allows us to specify which contemporaneous variables should appear in each equation.

While Krolzig (2003) used PcGets, we use the closely related search algorithm Autometrics, which is a package within the PcGive version 12 econometrics package (Doornik's chapter in this volume; Doornik and Hendry, 2007).

Since the data are all nonstationary, we reparameterize the SVAR with unrestricted lags in a vector-error-correction form:

$$\mathbf{A}_0 \Delta \mathbf{Y}_t = \mathbf{A}_1(L) \Delta \mathbf{Y}_{t-1} + \mathbf{A}_2 \mathbf{Y}_{t-1} + \mathbf{E}_t. \quad (2.3)$$

\mathbf{A}_0 is specified according to Figure 2.2 and is held fixed though all searches.

As $\Delta \mathbf{Y}_{t-1}$ is stationary, the estimates of the elements of \mathbf{A}_1 have standard distributions; while, since \mathbf{Y}_{t-1} is nonstationary, the estimates of the elements of \mathbf{A}_2 have nonstandard distributions and critical values need to be inflated in the direction of those provided by Dickey and Fuller (1979). PcGets does not allow the different critical values for different types of variables. Our strategy will be to conduct two ordinary-least-squares searches for each separate equation of the SVAR, using 5% and 1% critical values. Since these correspond to the default 'liberal' and 'conservative' settings that were built into PcGets, we continue to use that terminology. The critical values for these tests assume a normal distribution. The conservative setting is used here as an *ad hoc* method of mimicking the higher critical values of the nonstandard distributions appropriate for nonstationary variables. Each of our two searches has two stages:

- (i) A liberal search in which the contemporaneous variables and the lagged levels are held fixed; (ii) followed by a conservative search in which the contemporaneous variables and the short-term dynamics selected in stage (i) and indicated by the specification of \mathbf{A}_1 —the placement of zeros in the matrix, but not the estimates of the elements—are held fixed, while a search is conducted over the lagged level terms.
- (i) A conservative search in which only contemporaneous variables (identified by the nonzero elements of \mathbf{A}_0) and the lagged differences are held fixed; (ii) followed by a liberal search in which the contemporaneous variables and the long-term dynamics selected in stage (i) and indicated by the specification of \mathbf{A}_2 —the placement of zeros in the matrix, but not the estimates of the elements—are held fixed, while a search is conducted over the lagged differences.




The specification with the lowest Schwarz information criterion is chosen.

The detailed specification of the SVAR is estimated as a system by maximum likelihood. (While we do not report the detailed estimates here, they are available from the authors.)⁷ The complete causal order is summarized in Table 2.5. To get an idea of how much detail about the causal process our investigation has uncovered, how parsimonious our final specification is, consider what Table 2.5 would look like for any typical Choleski ordering: every cell on the main-diagonal and above would be filled in as light grey (indicating the presence of lagged variables) and every cell below the main diagonal in black (indicating the presence of contemporaneous and lagged

⁷ Estimates are available at <http://www.econ.duke.edu/~kdh9/research.html>.

Table 2.5. The causal structure of the SVAR

Effects	Causes										
	C O R E I N F	F F	I P	S P P E	M 2 O W N	L I Q D E P	S P S 0 0	R E F I	M O R G 3 0	V O L	T B I L L 3
COREINF											
FF											
IP											
SPPE											
M2OWN											
LIQDEP											
SPS00											
REFI											
MORG30											
VOL											
TBILL3											

Key:
 = lagged causes only
 = contemporaneous causes only
 = lagged and contemporaneous causes

variables). The overall lighter tone, especially the white (empty) cells, reveals that we have stripped away large numbers of redundant regressors.

2.6 The Role of M2 in the Monetary Transmission Process

We are now in a position to return to the business of the Monetary and Reserve Analysis Section. Table 2.5 is concerned with direct effects or proximate causes. The discussion in section 2.1 suggests that the entire *LIQDEP* row should be filled, indicating that each of the other variables is either a contemporaneous or lagged cause of *LIQDEP*. Most of the variables do cause *LIQDEP*; but the ones that do not are a surprise: neither the central quantity-theoretic variables, *COREINF* and *IP*, nor lagged *LIQDEP* itself causes *LIQDEP*. The 'special factors' are all selected as direct causes of *LIQDEP*. A pitfall of many statistical studies is to mistake causation for correlation. The pit is quite deep here, as *LIQDEP* is identified as a direct cause of each of the other variables except for *SPPE*, *M2OWN*, and *MORG30*.

Another pitfall of single-equation and correlational studies is the mixing of direct and indirect causes. Indirect causes can be investigated through impulse-response functions. In the interest of conserving space, we omit most impulse-response functions and focus on the ones most relevant to the role of M2 in the transmission of monetary policy.⁸

While most of the considerations concerning the relationship of M2 to other variables in section 2.1 are discussed from the point of view of the demand for money, the real importance of M2 to the Federal Reserve is found in its place in the transmission mechanism for monetary policy, in which M2 is a cause rather than an effect. Table 2.5 indicates that *LIQDEP* (and, therefore, M2) is a cause of both industrial production (*IP*), representing the real economy, and of core inflation (*COREINF*). But how important is this relationship in the transmission of monetary policy? This is difficult to quantify with the usual impulse-response functions, since there is no natural metric on which to compare the different effects of *LIQDEP*. In order to get a better handle on the role of *LIQDEP* in the transmission mechanism, we concentrate on the effects of a shock to the Federal funds rate (*FF*), the Federal Reserve's main policy variable. We consider a permanent, 25-basis-point shock, since this is by far the Federal Reserve's most common monetary policy action in practice.⁹ And since the Fed is principally concerned with the effects of its policies on growth and inflation, we focus on responses of core inflation (*COREINF*) and industrial production (*IP*) to the monetary-policy shock. Figure 2.3 displays the ordinary impulse-response functions, as well as impulse-response functions where *LIQDEP* has been shut down by setting its value to zero in all equations and all periods. The difference between these two impulse-response functions is a measure of the role of *LIQDEP* in the transmission mechanism.

The evidence of Figure 2.3 is that M2 (or *LIQDEP*) plays an almost immeasurably small role in the transmission mechanism. The effect of the increase in the Federal funds rate on core inflation (Figure 2.3) is not even statistically significantly different from zero except in period 4. The maximum difference between the impulse-response functions for *COREINF* occurs at period 3 and is less than a tenth of percentage point on the inflation rate. And the two impulse-response functions are nowhere statistically distinguishable in the sense that both lie within the standard-error bands for the function with *LIQDEP*. The differences between the two impulse-response functions for *IP* (Figure 2.3B) are similarly small, at most less than one-tenth of a percentage point on the growth rate of industrial production.

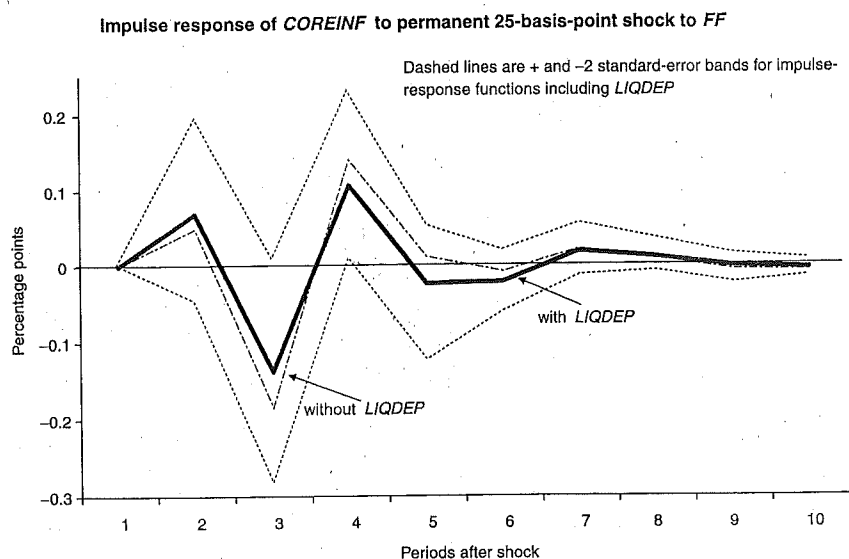
⁸ All of the impulse-responses of each of the other variables to *LIQDEP* and of *LIQDEP* to each of the other variables are available on Kevin Hoover's website (see footnote 7).

⁹ A permanent shock in the VECM form is a positive shock in the first period with no shocks in subsequent periods.

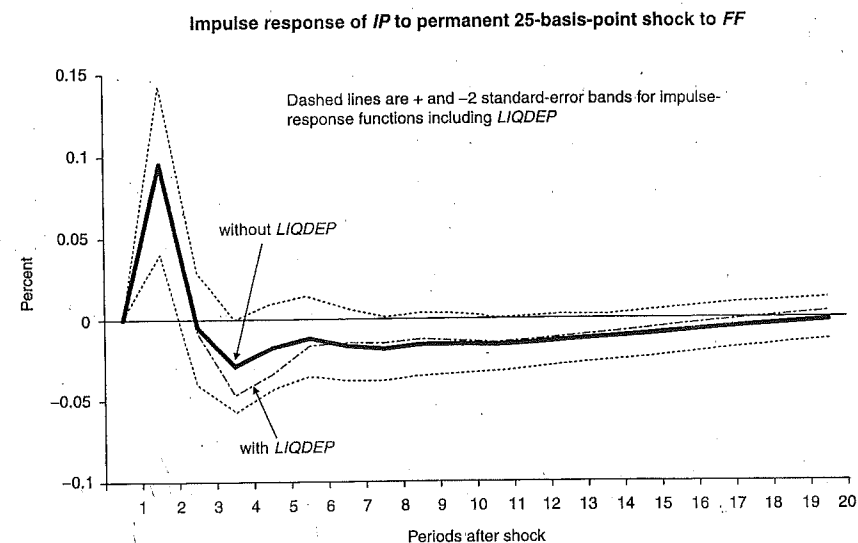
2.7 Conclusions

Our investigation is both methodological and substantial. Methodologically, it provides a concrete illustration of how to coordinate the graph-theoretic causal-search algorithms, previously applied to vector autoregressions by a number of investigators, with David Hendry's general-to-specific search methodology, embodied in PcGets and Autometrics, to identify empirically a structural econometric model in a case in which theory is relatively weak and not a reliable source of identifying restrictions. Typically, investigators use the PC algorithm or one of its relatives to select a contemporaneous causal graph. We showed how to use the bootstrap techniques developed by Demiralp, Hoover, and Perez to assess the uncertainties associated with selecting such a graph. These techniques proved invaluable guiding an informal general-to-specific search using test of overidentifying restrictions to select a contemporaneous causal order for the SVAR in which we could have reasonable confidence. One avenue for future development would be to provide a more formally developed search procedure along these lines. In applying Autometrics to nonstationary data, we adopted what we believe to be an effective, though *ad hoc*, procedure to ensure that appropriate selection criteria were applied to the I(1) as well as the I(0) terms. Another future development would be to extend Autometrics to do this automatically.

On the substantial side, we were able to provide a carefully tested fully identified model of the role of M2 in the transmission mechanism in which the informal assumptions of the Federal Reserve's Monetary and Reserve Analysis Section could be assessed. On the positive side, the evidence supports the section's identification of particular 'special factors' connected to the behaviour of M2. On the negative side, the quantity-theoretic core of their analysis seems to be largely at odds with the data. What is more, the rationale for focusing attention on M2 is undermined by evidence that it plays an insignificant role in the transmission of monetary policy. This is in keeping with other recent findings about M2 (Hale and Jordá, 2007). Some economists have argued that, while M2 has only a small role in the substantive transmission of monetary policy, it nonetheless conveys significant information about the state of the monetary economy in much the same way that a barometer, while not a causal linkage in weather systems, nonetheless provides useful information (see Dotsey and Otrok, 1994; for a survey of views). Our conclusions are different. Once we have such a detailed causal identification, we can see that M2 does play a real causal role (*LIQDEP* is identified as a cause of six of nine nonpolicy variables (other than itself) in our system). Its role is, nonetheless, a practically insignificant one. Our identified model conveys fuller information—and serves as a better barometer—than M2 can.



(A)



(B)

FIG. 2.3. The importance of *LIQDEP* in the transmission of monetary policy.

Appendix A. Data

All data are monthly. If a Haver series code is given, the data were downloaded from Haver Analytics *United States Economic Statistics* database, version date 3 June 2005. Natural logarithms indicated by 'log'.

$COREINF = 1200[\log(corecpi_t) - \log(corecpi_{t-1})]$, where *corecpi* = Consumer Price Index less Food and Energy. Source: Bureau of Labor Statistics [Haver: *pcuslfe*]. Units: Index 1982–84 = 100; seasonally adjusted.

FF = Effective Federal Funds Rate. Source: Federal Reserve [Haver: *ffed*]. Units: annual yield to maturity (percent per annum).

IP = $\log(\text{industrial production index})$. Source: Federal Reserve [Haver: *ip*]. Units: index 1997 = 100; seasonally adjusted.

LIQDEP = $\log(\text{liquid deposits})$, *liquid deposits* = Demand Deposits + Other Checkable Deposits + Savings Deposits. Source: Federal Reserve. Units: millions of dollars; seasonally adjusted.

M2OC = Opportunity Cost of Holding M2 = *TBILL3* – *M2OWN*. Source: Federal Reserve (Division of Monetary Affairs, Reserve Analysis Division).

M2OWN = Own Yield on M2 = weighted average of yields on deposits. Federal Reserve (Division of Monetary Affairs, Reserve Analysis Division); source for underlying yields: *Bank Rate Monitor*.

MORG30 = Interest Rate on 30-year Mortgages. Source: Federal Mortgage Acceptance Corporation, Primary Market Survey: 30-year Fixed Rate Mortgages [Haver: *frm30*]. Units: annual yield to maturity (percent per annum).

REFI = $\log(\text{mortgage refinancing})$. Source: Mortgage Banker's Association. Units: index 1990 = 100.

SVOL = $\log(\text{volatility})$, where *volatility* = VIX (up to August 2003), renamed VXO thereafter (both are measure of implied volatility based on the S&P 100 Index (OEX) option prices). Source: Chicago Board Options Exchange. Units: percent.

SP500 = $\log(\text{S\&P 500 stock index})$. Source: *Wall Street Journal* [Haver: *sp500*]. Units: index 1941–43 = 10.

SPPE = S&P 500 Price–Earnings Ratio. Source: Standard & Poors (<http://www2.standardandpoors.com/spf/xls/ondex/SP500EPSEST.XLS>) Units: ratio.

TBILL3 = Yield on 3-month Treasury bills. Source: Federal Reserve.

References

- Cooper, G. F. (1999). An overview of the representation and discovery of causal relationships using Bayesian networks. In Clark Glymour and Gregory F. Cooper (eds.), *Computation, Causation, and Discovery*, pp. 3–64. Menlo Park, CA and MIT Press, Cambridge, MA: American Association for Artificial Intelligence.
- Demiralp, S. and Hoover, K. D. (2003). Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and Statistics*, **65** (Supplement), 745–767.
- and Perez, S. J. (2008). A bootstrap method for identifying and evaluating a structural vector autoregression. *Oxford Bulletin of Economics and Statistics*, **70**, 509–533.
- Dickey, D. A. and Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, **74**(2), 427–431.
- Doornik, J. A. and Hendry, D. F. (2007). *PcGive Professional 12*, Vols. 1–3 and accompanying *PcGive*, version 12 software. London: Timberlake Consultants.
- Dotsey, M. and Otrok, C. (1994). M2 and monetary policy: A critical review of the recent debate. Federal Reserve Bank of Richmond, *Economic Quarterly*, **80**(1), 41–49.
- Gilbert, C. L. (1986). Professor Hendry's econometric methodology. *Oxford Bulletin of Economics and Statistics*, **48**(3), 283–307.
- Hale, G. and Jordá, O. (2007). Do monetary aggregates help forecast inflation? *Federal Reserve Bank of San Francisco Economic Letter* 2007–10, 13 April.
- Hendry, D. F. (1980). Econometrics: Alchemy or science? *Economica*, **47**(188), 387–406.
- and Krolzig, H.-M. (1999). Improving on 'Data Mining Reconsidered', by K. D. Hoover and S. J. Perez. *Econometrics Journal*, **2**, pp. 202–218.
- (2001). *Automatic Econometric Model Selection Using PcGets 1.0* and accompanying *PcGets*, version 1.0 software. London: Timberlake Consultants.
- Leamer, E. E. and Poirier, D. J. (1990). The *ET* dialogue: A conversation on econometric methodology. *Econometric Theory*, **6**(2), 171–261.
- Hoover, K. D. (2001). *Causality in Macroeconomics*. Cambridge: Cambridge University Press.
- (2005). Automatic inference of contemporaneous causal order of a system of equations. *Econometric Theory*, **21**(1), 69–77.
- and Perez, S. J. (1999). Data mining reconsidered: Encompassing and the general-to-specific approach to specification search. *Econometrics Journal*, **2**(2), 167–191.
- (2004). Truth and robustness in cross country growth regressions. *Oxford Bulletin of Economics and Statistics*, **66**(5), 765–798.
- Hurvich, C. M. and Tsai, C.-L. (1991). Bias of the corrected AIC criterion for underfitted regression and time series models. *Biometrika*, **86**(3), 499–509.
- Kant, I. (1787). *Kritik der reinen Vernunft*, 2nd edition.
- Krolzig, H.-M. (2003). General to specific model selection procedures for structural autoregressions. Working Paper, Department of Economics and Nuffield College Oxford, March.

- and Hendry, D. F. (2001). Computer automation of general-to-specific model selection procedures. *Journal of Economic Dynamics and Control*, **25**(6-7), 831-866.
- Mizon, G. E. (1995). Progressive modelling of economic time series: The LSE methodology. In Hoover, K. D. (ed.), *Macroeconometrics: Developments, Tensions and Prospects*, pp. 107-170. Boston: Kluwer.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Perez, S. J. (1998). Causal ordering and the 'bank lending channel'. *Journal of Applied Econometrics*, **13**(6), 613-626.
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, **48**(1), 1-48.
- Spirtes, P., Clark Glymour, C., and Scheines, R. (2001). *Causation, Prediction, and Search*, 2nd edition. Cambridge, MA: MIT Press.
- Scheines, R., Meek, C. and Glymour, C. (1994). *Tetrad II: Tools for Causal Modeling: User's Manual*.
- Swanson, N. R. and Granger, C. W. J. (1997). Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions. *Journal of the American Statistical Association*, **92**, 357-367.