ELSEVIER Journal of Banking & Finance 23 (1999) 1771–1792

Journal of BANKING & FINANCE

www.elsevier.com/locate/econbase

Does the stock market predict real activity? Time series evidence from the G-7 countries

Jongmoo Jay Choi [a,*], Shmuel Hauser [b,c], Kenneth J. Kopecky [a]

[a] *Department of Finance, Fox School of Business and Management, Temple University, Philadelphia, PA 19122, USA*

[b] *Ben Gurion University of the Negev, Beer Seva, Israel*

[c] *Israel Securities Authority, Jerusalem, Israel*

Received 19 March 1998; accepted 15 February 1999

Abstract

This paper extends one aspect of the US stock market study of Fama (1990) and Schwert (1990). We examine the relationship between industrial production (IP) growth rates and lagged real stock returns for the G-7 countries using both *in-sample* cointegration and error-correction models and the *out-of-sample* forecast-evaluation procedure of Ashley et al. (1980). The cointegration tests show a long-run equilibrium relationship between the log levels of IP and real stock prices, while the error-correction models indicate a correlation between IP growth and lagged real stock returns for all countries except Italy. The out-of-sample tests show that in several sub-periods the US, UK, Japanese, and Canadian stock markets enhance predictions of future IP. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: G12; G15

Keywords: Stock prices; Industrial production; G-7 countries

* Corresponding author. Tel.: +1-215-204-5084; fax: +1-215-204-1697.
E-mail address: v5156e@vm.temple.edu (J. Jay Choi).

PII: S 0 3 7 8 - 4 2 6 6 (9 9) 0 0 0 2 0 - 5

1. Introduction

The discounted-cash-flow valuation model states that stock prices reflect investors' expectations about future real economic variables such as corporate earnings, or its aggregate proxy, industrial production. If these expectations are correct on average, lagged stock returns should be correlated with the contemporaneous growth rate of industrial production. That is, real stock returns should provide information about the future evolution of industrial production. Fama (1990) examines this relation for the US over the 1953–87 period. In ordinary least-square (OLS) regressions, he reports an R^2 value of 0.14 for 12 monthly lags of real stock returns, 0.30 for 4 quarterly lags and 0.44 for a 1-year lag. Schwert (1990) expands the data set to include the entire 1889–1988 period for the US and finds results similar to those of Fama. However, as both authors admit, their results are subject to the limitations inherent in the use of an in-sample procedure such as OLS, which selects the explanatory variables on the basis of goodness-of-fit. Disentangling correlation and predictability within that framework is an impossible task.

The aim of our paper is to examine the relation between industrial production (IP) and lagged real stock returns using several different time series methodologies. Our goal is to build upon the work of Fama (1990) and Schwert (1990) in two directions. First, we use in-sample time series techniques to document the industrial production-stock return relation for both the US and the remaining G-7 countries. Since all of these countries have well-developed stock markets and high levels of output per capita, stock prices set by rational investors should exhibit patterns of correlation with the future growth of IP within each country. Second, we use an out-of-sample time series procedure developed by Ashley et al. (1980) (AGS) in order to avoid the potential problems associated with the use of in-sample procedures, such as OLS and Granger (1980) causality, in testing whether lagged stock returns predict industrial production growth. Our maintained hypothesis is the same as that in Fama (1990) and Schwert (1990). If the standard valuation model provides a correct empirical description of stock market behavior, lagged stock returns will contain forward-looking forecasts of IP and should, therefore, be useful in predicting the growth rate of industrial production. In our out-of-sample analysis, however, "usefulness" is not defined in an absolute sense, but rather relative to the past information that is already available in the history of industrial production. Of course, it is not necessary a priori for financial information to enhance real sector forecasts, because there is no guarantee that the financial sector is prescient with respect to future developments in the real sector. [1]

[1] In our empirical work we use IP as a proxy for aggregate corporate earnings because IP is the only aggregate data series available on a monthly basis. It should be noted that our findings may be sensitive to the choice of IP as the aggregate proxy variable.

However, a speedy adjustment of financial asset prices to *new* information suggests that access to financial information should enhance the empirical predictability of real economic activity relative to predictions based only on *dated* information about IP itself.

Other prior research in this area is focused on the interaction among the potential determinants of stock prices, a topic that is not the precise focus of this paper. Lee (1992) uses a standard VAR analysis to investigate the interactions among stock returns, real activity, interest rates, and the inflation rate in the US. His approach, however, is not based on cointegration analysis and thus may not have adequately addressed the potential unit root problem that is discussed below. Bittlingmayer (1992), Park (1997), and Gallinger (1994) examine the influence of various aspects of real activity on stock prices, which is the opposite relation to that examined in our paper. Closer in focus to our work is Canova and De Nicolo (1995), who use a three-country model comprised of the US, Europe, and the rest of the world to investigate linkages between domestic output growth and domestic stock returns, and Loflund and Nummelin (1997), who examine the interaction between asset returns and industrial production growth for Finland and Sweden.

Our results can be summarized as follows. Irrespective of whether monthly, quarterly or annual data are used, the in-sample cointegration analysis shows that the log levels of industrial production and real stock prices are characterized by a stationary linear relation in all G-7 countries. In addition, as determined by in-sample error-correction models estimated over all three data frequencies, real stock returns show significant evidence of short-run causality for the growth rate of industrial production in the US, UK, Japan, Canada, and Germany. In France, significant evidence of causality is found only at the quarterly frequency, while Italy fails to show causality at any data frequency. Thus, use of an in-sample time series methodology leads to results for the US that are consistent with those of Fama (1990) and Schwert (1990). Moreover, the Fama–Schwert findings appear to extend to five additional G-7 countries, in which IP growth is significantly correlated with lagged stock market information. The out-of-sample time series procedure, however, poses a different question and comes up with rather different results. For the out-of-sample forecasts, the key issue is whether lagged stock market information, which is assumed to contain forward-looking expectations, improves upon forecasts of IP growth that can be derived only from considering the past behavior of IP. The out-of-sample tests reveal that the relative value of stock market information depends importantly on both the periodicity of the data and the length of the in-sample estimation period relative to the length of the out-of-sample period. At a monthly frequency, we find evidence of an enhanced predictability of IP growth in Japan and the UK and perhaps the US. At a quarterly frequency, a significant improvement arises for the US, Canada and perhaps Germany.

It should be noted, however, that this support is not uniform across all of the estimated lag structures for real stock returns. [2]

Interpreted positively, one may conclude that the in-sample pattern of correlation between IP growth and lagged stock returns that is found in Fama (1990) and Schwert (1990) for the US can be extended to the UK, Canada, Germany, Japan, and France but not to Italy. When the question is posed in an out-of-sample time series framework, however, there are fewer instances in which the information available in stock market prices can be shown to provide significant additional insight into the future movements of IP. This finding may arise because IP growth is relatively easier to forecast on the basis of its own past so that stock market information is redundant; stock market expectations are too uninformed or too volatile to be of systematic assistance in forecasting future IP growth, or the variance of innovations in other determinants of stock prices (e.g., the risk premium and/or risk-free rate) is so high that it overwhelms the information value of real stock returns for IP growth. We discuss these possibilities later. We also discuss below some evidence bearing on the potential instability of the stock-return lag structure in the out-of-sample tests for the US. Such instability suggests that a rational forecaster may fail to take advantage of US stock market information on a systematic basis. That is, the forecaster may unwittingly choose an incorrect lag length and thus forego an improvement in the forecast accuracy of industrial production growth. In general, therefore, while the in-sample tests indicate the presence of correlation between IP growth and lagged stock returns, the out-of-sample results suggest that the ability to exploit this correlation is not pervasive throughout the G-7 countries.

The remainder of the paper is organized as follows: Section 2 describes the data used in the study. Section 3 presents the results of the in-sample cointegration and causality analysis. Section 4 describes the out-of-sample procedure, presents the associated test results, and offers an interpretation of the disparate findings. Section 5 contains a brief conclusion.

2. Data

The data consist of monthly observations of the aggregate stock price index, industrial production index, and consumer price index of the G-7 countries –

[2] In a more general setting, the relation between industrial production and lagged real stock returns may be influenced by additional variables such as the nominal interest rate, rate of inflation, and macroeconomic policy variables. As a bivariate method, the AGS procedure does not enable us to accommodate the potential joint effects of these other variables. Barro (1990), Kaul (1987) and Geske and Roll (1983), however, do examine such joint effects using in-sample regression methods.

Canada, France, Germany, Italy, Japan, the UK and the US. The data are from the *International Financial Statistics* of the International Monetary Fund. The sample period runs from January 1957 to March 1996.

Table 1 reports summary descriptive statistics of the sample. For the ease of comparison across the countries, the statistics are presented on a business cycle frequency for each G-7 country. Following Fama (1990) and Schwert (1990), we utilize industrial production to both measure real activity and define each country's business cycle. Depending on whether a monthly, quarterly or annual time period is used, real stock prices are obtained by dividing the period's nominal stock price index by the corresponding consumer price index. In our empirical work, Y_t refers to the percentage change (log difference) of industrial production over a given period and S_t is the similar percentage change of real stock prices. Table 1 presents the data on a business cycle frequency. The entries refer to the annualized percentage change (log difference) of industrial production and real stock prices calculated over each indicated growth or recession period. An overall comparison of average stock returns in growth periods to that in recession periods reveals that average stock returns in growth periods (8.18%) are higher than those in recession periods (4.77%). Thus, there is preliminary evidence regarding the positive association between economic growth and stock price movements in the G-7 countries.

3. In-sample cointegration and error-correction models

The time series analyses that are most complementary to the work of Fama (1990) and Schwert (1990) are in-sample cointegration and error-correction models. The first procedure investigates whether two non-stationary time series exhibit a stable linear relation, while the second can be used to examine the issue of in-sample causality of stock returns for industrial production.

It is necessary for cointegration that the individual time series be non-stationary. To investigate the properties of the data, we use the adjusted Dickey and Fuller (1981) test, according to which a time series x_t is non-stationary if $\beta = 1$ in the autoregressive representation,

$$x_t = \alpha_0 + \alpha_1 t + \beta x_{t-1} + \sum_{j=2}^{n} \gamma_j x_{t-j} + \varepsilon_t, \tag{1}$$

where x_t represents a time series, t is a time trend, ε_t is an error term, and α_0, α_1, β and γ_j are parameters. Since the underlying time series data are in levels, we first examine the stationarity of the log levels of industrial production and real stock prices. The first two columns under the unit root heading in Table 2 report the results of the adjusted Dickey–Fuller unit root tests at monthly, quarterly and annual frequencies for the G-7 countries. The test statistics

Table 1
Summary statistics of sample data[a]

Country	Growth periods				Recession periods			
	Business cycle		Annual % change of industrial production	Annual % real stock returns	Business cycle		Annual % change of industrial production	Annual % real stock returns
	From	To			From	To		
US	5/58	1/60	13.40	18.01	1/57	4/58	−9.63	−5.49
	3/61	8/69	6.87	5.07	2/60	2/61	−7.19	9.68
	8/71	11/73	8.48	1.88	9/69	8/71	−1.90	1.49
	4/75	12/78	9.00	6.14	12/73	3/75	−10.20	−8.90
	11/82	11/84	10.00	9.61	1/79	10/82	−2.40	7.92
	10/86	3/89	4.49	10.51	12/84	9/86	0.97	21.73
	6/93	2/95	5.87	6.51	4/89	5/93	0.35	9.77
					5/95	3/96	1.21	27.76
UK	3/59	5/60	9.40	27.14	1/57	2/59	−0.46	26.45
	2/63	2/65	8.77	5.07	6/60	1/63	−2.18	1.41
	9/67	6/69	6.24	13.79	3/65	8/67	0.72	3.54
	2/72	2/73	20.72	−6.42	7/69	1/72	−0.43	10.81
	9/75	6/79	5.39	17.87	3/73	8/75	−4.28	−12.95
	6/81	9/88	3.34	16.26	7/79	5/81	−7.37	13.24
	4/93	9/94	5.78	6.82	10/88	3/93	−0.61	8.99
					10/94	3/96	1.11	12.10
Japan	7/58	5/70	15.15	10.05	1/57	6/58	0.00	−2.68
	11/71	2/74	10.34	30.22	6/70	10/71	0.98	8.67
	4/75	2/80	7.76	8.04	3/74	3/75	−15.86	−2.50
	3/83	11/84	8.48	22.20	3/80	2/83	0.10	8.78
	6/87	11/90	6.39	−6.20	12/84	5/87	−0.05	4.24
	1/94	2/96	4.72	2.29	12/90	12/93	−4.46	−5.96
Germany	1/57	2/66	7.00	15.85	3/66	5/67	−8.14	−13.62
	6/67	11/70	8.57	5.54	12/70	12/71	−2.08	3.63
	1/72	8/73	6.90	0.72	9/73	7/75	−6.66	0.64
	8/75	2/80	4.50	0.39	3/80	8/84	−1.07	8.70
	9/84	6/91	3.49	10.24	7/91	12/95	−1.48	1.86
France	2/60	8/74	5.60	−0.63	1/57	1/60	4.04	15.26
	9/75	8/79	5.55	8.92	9/74	8/75	−9.02	4.25
	1/87	8/90	4.02	5.93	9/79	12/86	−0.26	16.61
	1/94	8/95	4.38	−12.08	9/90	12/93	−1.09	13.39
					9/95	12/95	−5.79	14.69
Italy	5/58	1/64	11.02	11.08	1/57	4/58	2.51	−0.85
	9/65	8/69	8.88	5.22	2/64	8/65	0.42	−7.22
	2/73	6/74	16.63	10.76	9/69	1/73	2.64	−8.74
	10/75	3/80	6.69	7.65	7/74	9/75	−8.72	−24.43
	8/83	12/89	4.20	21.38	4/80	7/83	−2.64	28.63
					1/90	12/93	−1.00	−4.17

Table 1 (Continued)

Country	Growth periods				Recession periods			
	Business cycle		Annual % change of industrial production	Annual % real stock returns	Business cycle		Annual % change of industrial production	Annual % real stock returns
	From	To			From	To		
Canada	1/59	3/60	7.29	−6.89	1/57	12/58	1.25	−1.82
	5/61	3/69	7.47	6.58	4/60	4/61	2.81	23.64
	11/70	5/74	8.04	3.18	4/69	10/70	0.11	−13.06
	11/75	9/79	4.51	15.93	6/74	10/75	−5.52	−5.57
	1/83	5/88	6.25	9.06	10/79	12/82	−5.00	6.79
	1/92	1/95	5.49	3.65	6/88	12/91	−2.59	0.58
					2/95	3/96	−0.63	17.32
Average			7.68	8.18			−2.43	4.77
S.d.			3.65	8.88			4.17	11.75
Median			6.89	7.23			−1.07	4.24

[a] The duration of the business cycle is determined by the turning points of the industrial production index. The growth rate of real activity over each phase of the cycle is measured by the annualized monthly percentage change of the industrial production index. Similar business cycle calculations are made to obtain the annualized percentage change in real stock prices.

indicate that the hypothesis of non-stationarity cannot be rejected in any of the G-7 countries for either time series, irrespective of the time frequency considered. Thus, the time series in log levels satisfy the necessary condition for cointegration. We next examine the stationarity of the first difference of each time series (i.e., the log difference of real stock prices and industrial production), and, as shown under the first difference columns in Table 2, find that the null hypothesis of non-stationarity is rejected for both series. IP growth and real stock returns, therefore, cannot be cointegrated.

Given the non-stationarity of the log levels of IP and real stock prices, the cointegration of these two time series is addressed by estimating the following regressions:

$$y_t = \alpha + \beta s_t + \varepsilon_t, \tag{2}$$

$$\varepsilon_t - \varepsilon_{t-1} = -b\varepsilon_{t-1} + \theta_t, \tag{3}$$

where y_t and s_t represent the log levels of industrial production and real stock prices, respectively, and ε_t and θ_t are error terms. Following Engle and Granger (1987), if b is significantly different from zero, y_t and s_t are said to be cointegrated. The results of this significance test are given under the cointegration column in Table 2 and show that the log levels of IP and real stock prices are

Table 2
In-sample unit root, cointegration and causality tests[a]

Country		Unit root tests				Cointegration	Causality-ECM		
		Level		First difference		Stock and IP	χ^2	p-value	R^2
		Stock	IP	Stock	IP				
US	M	−1.35	−2.37	−8.48*	−8.39*	84.6*	2.90	0.000	0.36
	Q	−1.23	−2.85	−5.78*	−6.21*	33.1*	8.75	0.000	0.53
	Y	−0.92	−2.46	−6.12*	−5.30*	8.0*	4.05	0.052	0.84
UK	M	−2.07	−2.66	−9.29*	−9.87*	72.7*	2.11	0.016	0.18
	Q	−1.80	−3.09	−5.33*	−5.91*	22.3*	4.29	0.002	0.28
	Y	−1.46	−2.60	−4.82*	−4.72*	3.63**	9.23	0.000	0.83
Japan	M	−1.86	−0.64	−9.10*	−6.74*	74.9*	1.95	0.027	0.44
	Q	−2.08	−2.06	−5.10*	−7.30*	29.6*	4.22	0.003	0.56
	Y	−2.17	−2.49	−4.49*	−5.44*	7.3*	7.78	0.002	0.69
Canada	M	−3.06	−1.30	−8.41*	−7.59*	49.3*	2.61	0.000	0.24
	Q	−2.93	−1.72	−6.17*	−5.80*	16.4*	8.36	0.000	0.30
	Y	−2.70	−1.72	−5.58*	−3.99*	4.1*	4.22	0.022	0.86
France	M	−1.69	−1.10	8.88*	−10.27*	56.7*	1.09	0.365	0.26
	Q	−1.55	−1.01	−5.80*	−6.83*	21.6*	3.21	0.015	0.26
	Y	−1.16	−1.73	−5.37*	−4.72*	7.6*	0.86	0.439	0.59

Germany	M	−3.12	−2.13	−8.34*	−11.06*	80.2*	2.84	0.000	0.38
	Q	−3.10	−2.40	−4.81*	−5.62*	19.3*	4.87	0.001	0.35
	Y	−3.07	−2.95	−5.27*	−5.02*	8.1*	8.99	0.005	0.71
Italy	M	−1.93	−2.13	−7.92*	−10.26*	46.0*	1.08	0.411	0.29
	Q	−2.50	−2.26	−4.97*	−6.31*	14.3*	1.66	0.164	0.24
	Y	−1.76	−1.93	−4.25*	−5.01*	19.4*	1.36	0.276	0.60

[a] Unit root tests are based on $x_t = \alpha_0 + \alpha_1 t + \beta x_{t-1} + \sum_{j=1}^{n} \gamma_j x_{t-j} + \varepsilon_t$, where x_t is defined as either the log of industrial production, log of real stock prices, growth rate of industrial production, or real stock returns. The null hypothesis of non-stationarity, $H_0 : \beta = 1$ is tested using the adjusted Dickey and Fuller (1981) statistic. Cointegration tests are based on $y_t = \alpha + \beta s_t + \varepsilon_t;\ \varepsilon_t - \varepsilon_{t-1} = -b\varepsilon_{t-1} + \theta_t$, where Y_t and S_t represent the log levels of industrial production and real stock prices, respectively. The null hypothesis of no cointegration is $H_0 : b = 0$. The error correction model (ECM) is $Y_t = \sum_{j=1}^{m} a_{t-j} S_{t-j} + \sum_{i=1}^{m} b_j Y_{t-i} + \delta\varepsilon_{t-1} + \phi_t$, where Y_t and S_t represent the growth rate of industrial production and real stock returns, respectively, ε_{t-1} is the residual from the cointegration equation and ϕ_t is an error term. The causality test is based on the statistic $([\mathrm{SSE}(Y) - \mathrm{SSE}(Y,S)]/v)/(\mathrm{SSE}(Y,S)/(n-2v-1))$, which is distributed as χ^2 with v degrees of freedom and n observations. $\mathrm{SSE}(Y,S)$ is the sum of squared residuals of the error correction model, and $\mathrm{SSE}(Y)$ is the sum of squared residuals from the reduced model, $Y_t = b_0 + \sum_{i=1}^{m} b_j Y_{t-i} + \eta_t$, where η_t is an error term. The P-value column reports the level of significance for the null hypothesis of no causality. R^2 refers to the ECM model. M, Q and Y are tests based on monthly, quarterly and annual data, respectively.

* Indicates significance at 1% level.

** Indicates significance at 5% level.

cointegrated in each of the G-7 countries at either monthly, quarterly or yearly frequencies. Thus, from a long-run equilibrium perspective neither of these two variables drifts apart in any of the G-7 countries.

Given that the log levels of IP and real stock prices are cointegrated, we estimate an error correction model for the growth rate of industrial production. This model expresses the growth rate of industrial production (Y_t) in terms of lagged IP growth rates, lagged real stock returns (S_t) and a lagged cointegration error term:

$$Y_t = \sum_{j=1}^{m} a_{t-j} S_{t-j} + \sum_{i=1}^{m} b_j Y_{t-i} + \delta \varepsilon_{t-1} + \phi_t, \tag{4}$$

where ε_{t-1} is the residual from Eq. (3) and ϕ_t is a random error term.

Eq. (4) bears the closest resemblance to the specification used by Fama (1990) and Schwert (1990). However, the error correction term, $\delta\varepsilon_{t-1}$, is an additional element that influences the estimated coefficients in the lags of IP growth and real stock returns. These lagged coefficients are not unconstrained but must keep the underlying log-level series sufficiently close to their long-run equilibrium relationship. The results of the error-correction model are reported under the causality-ECM heading in Table 2. We first note that the R^2 value for the US. at any frequency is higher than the comparable R^2 reported by Fama. This arises because estimating Eq. (4) without an error-correction term represents a misspecification when the log-levels of the two series are cointegrated. More importantly, the results indicate that for the three data frequencies considered five countries, the US, UK, Japan, Germany, and Canada, show significant short-run in-sample causality of real stock returns for IP growth at approximately the 5% level or better in a two-tailed χ^2 test. [3] There is also evidence of causality in France at the quarterly frequency. Only Italy fails to show any support for the in-sample causality of real stock returns.

The in-sample cointegration and error-correction models provide both a broad confirmation of the Fama–Schwert finding for the US and an extension of their result to five of the six remaining G-7 countries, with Italy as the sole exception. Our results are also consistent with those of Kaul (1987), who uses a regression approach similar to that of Fama (1990) in modeling IP growth in Canada, Germany, the UK and the US and reports a significant influence of lagged stock returns on IP growth in each country. The relation between IP growth and lagged stock returns has, therefore, been examined with different

[3] The values of the lag length, m, are 24, 8 and 2 for monthly, quarterly, and annual data, respectively.

estimation techniques, over different time periods, and with different countries included. The conclusion seems inescapable. When the investigator possesses knowledge of the data over the entire sample period and uses that data to estimate the in-sample relationship between IP growth and real stock returns, lagged values of real stock returns are found to be significantly correlated with the contemporaneous value of the IP growth rate in six of the G-7 countries.

4. Out-of-sample AGS prediction tests

4.1. Methodology

The in-sample analysis shows that IP growth and lagged stock returns are correlated over the entire sample period in six out of seven G-7 countries. An in-sample correlation, however, is an ex post property of the data. An alternative analysis focuses on the out-of-sample predictive power of time series models and thus provides an ex ante view of the causal relation between IP growth and lagged real stock returns. According to this approach, the importance of stock returns for industrial production is determined by its ability to improve forecasts of industrial production over a portion of the sample that has *not* been used to estimate the underlying relationships.

For the out-of-sample tests, we use a minor modification of a procedure developed by Ashley et al. (1980). Letting Y represent the rate of growth of industrial production and S real stock returns, AGS define causality in the time series sense as follows:

> S is said to *cause* Y if $\mathrm{MSE}(Y,S)$ is significantly smaller than $\mathrm{MSE}(Y)$, where $\mathrm{MSE}(Y)$ is the population mean-square of the one-step-ahead forecasting errors of Y_{t+1} using a linear model based on Y_{tj}, and $\mathrm{MSE}(Y,S)$ is the population mean-square of the one-step-ahead forecasting errors of Y_{t+1} using a bivariate linear model based on Y_{tj} and S_{tj}, where $j = 0, 1, \ldots, m$ and signifies the number of lags.

AGS use an ARIMA model to determine the structure of the linear forecasting equation from which $\mathrm{MSE}(Y)$ is obtained, and a heuristic cross-correlogram analysis to determine the structure of the linear bivariate forecasting equation that is used to obtain $\mathrm{MSE}(Y,S)$. Because of specification and sampling errors, no straightforward test is available to determine the significance of the relative improvement in the one-step-ahead forecasts.

To test for significant differences between $\mathrm{MSE}(Y)$ and $\mathrm{MSE}(Y,S)$, AGS propose the following test procedure. Denote e_1 and e_2 as the forecasting errors of the univariate Y-model and the bivariate Y, S-model, respectively and define

$d_t = e_{1t} - e_{2t}$ and $\Sigma_t = e_{1t} + e_{2t}$. Letting COV and μ represent the out-of-sample covariance and mean, respectively, AGS show that the mean-square errors of the univariate model are significantly greater than the mean-square errors of the bivariate model, if the null hypothesis of $\mathrm{COV}(d, \Sigma) = 0$ and $\mu(d) = 0$ is rejected in favor of the alternative hypothesis that both of these quantities are non-negative and at least one is positive. AGS show that this is equivalent to testing the null hypothesis that $a_0 = a_1 = 0$ against the alternative that both are non-negative and at least one is positive in the following regression: [4]

$$d_t = a_0 + a_1[\Sigma_t - \mu(\Sigma_t)] + w_t, \tag{5}$$

where w_t represents a random error term.

To implement Eq. (5), we use the following procedure: (a) an in-sample univariate ARIMA model is estimated for each time series of the growth rate of industrial production using the Box and Jenkins (1970) procedure; (b) an in-sample bivariate model of the growth rate of industrial production that includes both lagged industrial production growth and lagged real stock returns is also estimated; (c) the univariate and bivariate models are used to generate one-step-ahead forecasts for a *preassigned* out-of-sample period; and (d) the AGS test based on Eq. (5) is used to compare the relative significance of the forecasting errors of the two models. [5]

Following the AGS (Ashley et al., 1980) procedure, we estimate the following in-sample univariate and bivariate models by ordinary least-squares to obtain our out-of-sample forecasting equations:

$$Y_t = a_0 + \sum_{i=1}^{m} a_i Y_{t-i} + e_{1t}, \tag{6}$$

$$Y_t = b_0 + \sum_{i=1}^{m} b_i Y_{t-i} + \sum_{j=1}^{n} c_j S_{t-j} + e_{2t}. \tag{7}$$

[4] For a detailed description of this test, see AGS (Ashley et al., 1980, pp. 1154–1155).

[5] The second step (b) in the above procedure requires the identification and estimation of an in-sample bivariate model that is to be used for out-of-sample forecasting. In AGS paper (Ashley et al., 1980), both economic theory and heuristic in-sample information about the cross-correlograms of the residual ARIMA series are used to determine a specific set of lags for the bivariate model. For our study, however, economic theory is vague about the precise lag relationship between industrial production growth and real stock returns. In the absence of economic theory and to avoid the use of heuristic information, we use in-sample information to estimate an upper bound on the length of the lag structure for real stock returns in the bivariate model of industrial production growth. It should be noted that our focus on the length of the lag structure rather than on the choice of a specific set of lags represents a minor modification of the original AGS (Ashley et al., 1980) procedure.

The length of the lag for industrial production growth, m, is set equal to 3 in the model using monthly data and to 2 in the quarterly model. The lag length for real stock returns, n, is set equal to 12, 18 or 24 in the monthly model and to 4, 6 or 8 in the quarterly model. For the monthly model, the $m = 3$ lag in Eq. (6) is generally consistent with the results of the ARIMA models in which the length of the autoregressive lag for industrial production never exceeds three months, irrespective of the assumed length of the in-sample period. For the quarterly model, a lag length of 2 is also generally consistent with the ARIMA results. Regarding the length of the lag for real stock returns, n, we use a χ^2 test to search for a significant lag length and find that it differs in each country. Therefore, we report the results of the AGS test in Eq. (7) for each of the alternative monthly and quarterly lag lengths of n.

We derive results for two arbitrary out-of-sample forecasting periods. The first uses 50% of the data for each country to construct in-sample estimates of Eqs. (6) and (7) and thus the remaining 50% is used to evaluate the forecasting properties of Eq. (7) relative to Eq. (6). The 50% in-sample period runs from January 1957 to July 1976. The second out-of-sample period uses 75% of the data for in-sample estimation and 25% for out-of-sample forecast evaluation. The 75% in-sample period runs from January 1957 to April 1986. AGS (Ashley et al., 1980) use an 80/20 in/out ratio in their empirical work but suggest that the out-of-sample period should be lengthened to provide more accurate out-of-sample tests. Our choice of the two out-of-sample periods follows the AGS suggestion. [6]

Following the AGS procedure, we estimate ARIMA models for industrial production growth in each of the G-7 countries based on in-sample data. These models are estimated for the cases when the in-sample period contains either 75% or 50% of the available data. The results are reported in Table 3 for both monthly and quarterly frequencies. Note, in addition, that the Ljung–Box test for white noise residuals is performed for up to either 12 or 24 lagged residuals.

Using monthly data, the autoregressive lag structure in four countries, Canada, Italy, France and the UK, is unaffected by the assumed length of the in-sample period. In addition, the hypothesis of white noise residuals cannot be rejected at the 5% level in any of these four countries. On the other hand, in the ARIMA models for the US, Germany and Japan, the reported χ^2 values in Table 3 indicate rejection of the null hypothesis of white noise residuals at the 5% level for either the 12 or 24 lagged residual length. However, at the more stringent 1% level the white noise hypothesis cannot be rejected in these three countries. Overall, therefore, the monthly growth rate of industrial production

[6] It should be emphasized that for the purpose of generating one-step-ahead out-of-sample forecasts, the univariate and bivariate models are re-estimated on the basis of new information.

Table 3
In-sample ARIMA models of industrial production growth rates: A calibration for the AGS model[a]

Country	In-sample 50%					In-sample 75%				
	p	d	q	χ^2_{12-p-q}	χ^2_{24-p-q}	p	d	q	χ^2_{12-p-q}	χ^2_{24-p-q}
Monthly data										
US	2	0	2	16.3*	33.8*	3	0	2	9.2	33.1*
UK	1	0	2	7.8	26.3	2	0	1	9.3	31.9
Japan	3	0	2	14.5*	31.3*	3	0	3	13.5*	28.1
Canada	2	0	2	3.9	25.0	2	0	2	14.9	22.9
France	2	0	0	5.3	7.5	2	0	1	2.3	4.8
Germany	2	0	4	19.7*	28.3	2	0	4	14.2*	26.8
Italy	2	0	2	9.9	27.9	2	0	2	6.8	25.7
Quarterly data										
US	2	0	2	7.5	11.6	2	0	2	8.5	20.0
UK	1	0	2	6.4	13.6	2	0	2	9.7	15.8
Japan	3	0	2	13.0	30.5*	2	0	2	18.4*	33.9*
Canada	2	0	2	10.1	20.7	2	0	2	9.1	17.6
France	2	0	0	12.8	15.9	2	0	0	18.8*	25.1
Germany	2	0	2	12.5	17.0	2	0	2	13.1	7.5
Italy	2	0	2	9.6	21.4	2	0	2	9.7	17.1

[a] The table indicates the number of significant terms for the autoregressive (p), moving average (q) and difference (d) components in the ARIMA models of the logarithmic difference of industrial production for the alternative countries, where the in-sample length refers to the use of either 50% or 75% of the entire data set. Results are reported using either monthly or quarterly data. The Ljung–Box test for model adequacy is

$$\chi^2 = N\sum_{k=1}^{L} r_k^2(\varepsilon),$$

where ε represents the moving average shock terms, r_k is the residual autocorrelation of lag k, L is the number of residual autocorrelations (either 12 or 24 in our tests), and N is the number of residuals in the estimated equations (number of observations minus L). χ^2 is distributed χ^2 with $L - p - q$ degrees of freedom. The null hypothesis assumes white noise residuals.
* Represents rejection of the null at the 5% level.

is adequately represented by an ARIMA process in four of the relatively smaller G-7 countries, while it is only marginally acceptable in the three largest countries.

Using quarterly data, Japan is the only country that systematically rejects the hypothesis of white noise residuals at the 5% level, but not at the more stringent 1% level. There is also a marginal rejection in France for the 75% in-sample period at 12 lagged residuals. When the residual lag is increased to 24, however, the hypothesis of white noise residuals cannot be rejected in France. Overall, therefore, ARIMA modeling of IP is satisfactory in six of the G-7 countries. Only Japan indicates marginal acceptance of the ARIMA model for quarterly data.

4.2. AGS out-of-sample tests

Table 4 displays the results of the out-of-sample tests of the relative importance of real stock returns for predictions of IP growth for two out-of-sample forecasting periods: a 50/50 in/out-of-sample division of the data and a 75/25 in/out division. Because the lag structure of real stock returns in the bivariate model is indeterminate theoretically, we examine three alternative lag lengths: 12, 18, and 24 months in the monthly model and 4, 6, and 8 quarters in the quarterly model. The null hypothesis of a shorter real-stock-return lag length can be tested at both 18 and 24 monthly lags and at 6 and 8 quarterly lags. The column labeled p_{12} provides the p-value for the null hypothesis of 12 monthly lags against the alternative of 18 monthly lags, p_{18} for the null of 18 lags against 24 monthly lags, p_4 for the null of 4 versus 6 quarters, and p_6 for the null of 6 against 8 quarters.

We first examine the monthly results in Table 4. The shape of the in-sample lag length for real stock returns in the monthly model varies by both country and the length of the in-sample estimation period. In the case of the US, the p-values for the 50/50 division indicate that one would reject 24 lags in favor of 18 and then reject 18 lags in favor of 12. On the other hand, for the 75/25 division, one would continue to reject 24 lags in favor of 18 but accept 18 lags rather than 12. Thus, increasing the length of the in-sample estimation period lengthens the lag on real stock returns for the US. Similar logic indicates a lag length of 24 months for the UK and Japan, and 12 months for Canada and Germany. France and Italy behave in a fashion similar to that of the US. In the case of France, the 50/50 data division implies an 18-month lag, while the 75/25 division leads to a 24-month lag. For Italy, the lag is 12 months for the 50/50 division and 18 months for the 75/25 division.

In the AGS procedure, causality is defined in terms of the relative improvement in the out-of-sample one-step-ahead forecasts. [7] Inspection of the monthly panel in Table 4 shows that the relative impact of using lagged real stock returns to predict the real sector is quite diverse across the G-7 countries. First, there is no confirming evidence of a significant stock market effect on monthly forecasts of IP growth in Canada, France, Germany, and Italy. As

[7] We have also performed out-of-sample tests for the case in which the in-sample period is truncated 12 months prior to the last observed business cycle turning point in each country. Out-of-sample forecasting was conducted for at least 24 months, 12 months prior to the last turning point and a minimum of 12 months following the last turning point. None of the out-of-sample test statistics were significant. However, two aspects of this experiment may lead to biased results. First, using a single turning point for out-of-sample forecasting in each country may create a sample size problem. Second, knowledge of the dating of a turning point implies overly strong a priori information on the part of the econometrician. The test statistics in Table 4 should be immune to both of these problems.

Table 4
AGS tests of lagged real stock returns for industrial production growth: Alternative out-of-sample forecasting periods and lag lengths for real stock returns[a]

Monthly data						
Country	In/out-of-sample period	12 lags	18 lags	p_{12}	24 lags	p_{18}
US	50/50	0.947	0.874	0.851	0.897	0.220
	75/25	1.055	1.079	0.019	1.092**	0.468
UK	50/50	1.068	1.049*	0.093	1.338*	0.021
	75/25	1.152	1.089*	0.030	1.070*	0.078
Japan	50/50	1.162*	1.142*	0.027	1.139*	0.072
	75/25	1.211	1.139*	0.004	1.150*	0.081
Canada	50/50	0.874	0.928	0.273	0.953	0.690
	75/25	1.036	1.202	0.169	1.192	0.602
France	50/50	0.521	0.516	0.001	0.514	0.501
	75/25	0.525	0.909	0.001	0.859	0.033
Germany	50/50	0.865	0.864	0.689	0.812	0.220
	75/25	0.984	1.074	0.615	1.086	0.168
Italy	50/50	0.678	0.592	0.200	0.582	0.110
	75/25	0.981	0.976	0.098	0.983	0.384
Quarterly data						
Country	In/out-of-sample period	4 lags	6 lags	p_4	8 lags	p_6
US	50/50	1.339*	1.441*	0.780	1.363*	0.497
	75/25	0.847	1.065	0.986	1.092	0.998
UK	50/50	1.061	1.127	0.839	1.220	0.529
	75/25	0.809	1.056	0.320	1.374	0.805
Japan	50/50	1.001	1.044	0.976	1.118	0.223
	75/25	1.035	1.058	0.783	1.153	0.68
Canada	50/50	1.256*	1.288*	0.548	1.378*	0.246
	75/25	1.050	1.107	0.891	1.130	0.318
France	50/50	0.825	0.850	0.783	0.963	0.869
	75/25	0.749	0.768	0.360	1.220	0.338
Germany	50/50	1.105	1.204	0.465	1.315*	0.305
	75/25	0.949	1.112	0.415	1.314	0.417
Italy	50/50	0.992	1.185	0.653	1.640	0.679
	75/25	1.089	1.384	0.271	1.802	0.288

[a] The table reports the ratio of MSE(Y) – the mean-square forecast error of the univariate model of industrial production growth – and MSE(Y,S), the mean-square forecast error of the bivariate model of industrial production growth and real stock returns. This test has been suggested by AGS (Ashley et al., 1980), in which the following regression is estimated:

$$\delta_t = \alpha_0 + \alpha_1[\Sigma_t - \mu(\Sigma_t)] + W_t,$$

where $\delta_t = \varepsilon_{1t} - \varepsilon_{2t}$, $\Sigma_t = \varepsilon_{1t} + \varepsilon_{2t}$, and ε_1 and ε_2 are the forecasting errors of the univariate and bivariate models, respectively. If $\alpha_0 = \alpha_1 = 0$, the null hypothesis cannot be rejected against the alternative that both are non-negative and at least one is positive. Out-of-sample forecasting is conducted over two alternative divisions of the data, one containing a 50/50 split between estimation and forecasting periods and the other a 75/25 split. The p-values refer to tests of the significance of alternative lag lengths for real stock returns in the bivariate model of industrial

Table 4 (Continued)
production growth. The null for p_{12} is 12 monthly lags of real stock returns against 18 lags, while the null for p_{18} is 18 lags against the alternative of 24 monthly lags. The null for p_4 is 4 quarterly lags versus 6, and for p_6 the null is 6 quarterly lags against the alternative of 8.
* Indicates that this ratio is significantly greater than unity, implying that the null hypothesis of no causality is rejected at the 5% level.
** Indicates that this ratio is significantly greater than unity, implying that the null hypothesis of no causality is rejected at the 10% level.

noted above, the shortest feasible lag length of real stock returns for Canada and Germany is 12 months. Thus, the findings against a significant out-of-sample stock market effect in these two countries may simply reflect the fact that a significant in-sample lag for real stock returns, if any, is shorter than 12 months. We have not investigated a shorter lag length because of the need to capture at a minimum a full seasonal (i.e., 12 months) pattern of lagged effects of real stock returns. A similar rationale for using 12 monthly lags underlies the Fama–Schwert methodology.

Second, the UK and Japan provide striking evidence in support of an enhanced real sector predictability at the 5% level, irrespective of the length of the in-sample estimation period. In both countries, selection of a 24-month stock-return lag, which is consistent with the p-values in Table 4, provides improved forecasts of IP growth relative to the forecasts made *without* using stock return data. Moreover, the use of an even shorter, and potentially inferior, 18-month lag length also produces significant evidence of stock market causality. Thus, in the case of both the UK and Japan, a rational forecaster, using hypothesis tests to determine the lag length of real stock returns in the estimated in-sample equations, would have been able to improve the out-of-sample forecasts of IP growth relative to alternative forecasts based only on the history of IP growth.

How does one interpret the monthly results for the US in Table 4? The table indicates that lagged US stock returns communicate relatively valuable information about the monthly movements of US industrial production at the 10% level for the 75/25 division of the data. However, this positive evidence occurs at 24 monthly lags of real stock returns rather than at the significant in-sample lag length of 18 months. Based on in-sample estimation, therefore, it is likely that a rational forecaster would have failed to exploit the in-sample information that is evidently contained in the 24-month lag.

Monthly data are inherently more volatile than quarterly data. Thus, the question arises whether the quarterly out-of-sample predictions of IP growth can be improved through the use of lagged quarterly real stock returns. The panel referring to quarterly data in Table 4 presents evidence on this issue. First, and in sharp contrast to the monthly results, the p-values indicate that the in-sample lag length is uniformly 4 quarters, irrespective of either the specific country considered or the arbitrary division of the data into in-and out-of-sample periods. Second, in the UK, Japan, France, and Italy lagged stock returns have no significant additional effect on the predictions of quarterly IP growth. This result concurs with the monthly findings for France and Italy, but is at odds with the monthly results for the UK and Japan, in which improved IP growth forecasts are achieved with stock return lag lengths of either 18 or 24 months. Third, the German stock market appears to enhance forecasting of the real sector when 8 quarterly lags are used for in-sample estimation. However, given the available in-sample p-values in Table 4, it is unlikely that a rational forecaster would have selected 8-quarter lag length for real stock returns.

For the US and Canada, the quarterly results, in contrast to the monthly findings, reveal an entirely different picture regarding the relative importance of lagged stock returns. Irrespective of the length of the lag structure, stock returns lead to improved forecasts of IP growth at the 10% level for the 50/50 data division. Most importantly, this significant enhancement is consistent with the a priori choice of lag structure by a rational forecaster, since 4 quarterly lags would have been selected on the basis of in-sample information. As the data division is changed to the 75/25 split, however, the significant influence of lagged stock returns on IP growth forecasts disappears in both countries. This result shows that the assumed length of the in-sample estimation period relative to the out-of-sample forecasting period can importantly influence the inferences drawn under the AGS procedure. As noted above, AGS (Ashley et al., 1980) argue that a 20% out-of-sample forecasting period is likely to be too short, which supports our choice of longer 25% and 50% out-of-sample periods.

4.3. Interpretation of the AGS results

How can we reconcile the different findings between the in-sample error-correction models and the out-of-sample AGS tests and between the monthly and quarterly AGS tests? The in-sample time series estimation procedure is designed to find patterns of correlation between IP growth and lagged real stock returns. These in-sample correlation patterns can emerge, however, without any implication that stock returns are prescient about the future path of industrial production. For example, suppose that real stock returns and IP growth are contemporaneously correlated. If IP growth exhibits persistence, in-sample estimation will discover that lagged stock returns are correlated with current IP growth. On the other hand, the AGS out-of-sample procedure seeks to determine whether the estimated in-sample correlation patterns can provide an advantage to a forecaster who has access to an alternative simpler forecasting instrument, namely, a model that uses only lags of IP growth to forecast itself. In this view, IP evolves over time independently of stock returns and is determined by real sector forces such as technological change and other labor force and demographic characteristics. These influences change gradually and thus IP growth exhibits sluggish, backward-looking behavior. Since the value of the aggregate stock market is assumed to depend on the future evolution of IP, forward-looking investors have an incentive to anticipate innovations in the path of IP, which in turn would be incorporated into immediate changes in real stock returns. Assuming that these revised expectations are on average correct, a forecaster, using data on real stock returns, would thus be able to anticipate a component of future IP growth that could never be predicted from the past history of IP itself. In this restricted sense, therefore, the AGS procedure allows us to determine positive instances of prescience in the stock market, an issue that cannot be addressed by any in-sample estimation procedure. The disparate findings between the error correction models and the AGS tests are not in conflict because these two time series approaches are not investigating the same question. IP growth and lagged stock returns may exhibit a high degree of in-sample correlation, but lagged stock returns may, nonetheless, perform relatively poorly in forecasting out-of-sample IP growth.

To interpret the international differences in the AGS results, we should first realize that lagged stock returns may be especially advantageous in forecasting IP growth whenever there is a potential difficulty in modeling IP growth with its own lags. The ARIMA results in Table 3 can be used to examine this issue. In the monthly panel, there are eight starred values under the χ^2 columns that indicate some degree of uncertainty about the overall adequacy of the ARIMA model for IP growth. The AGS test results in Table 4 confirm the relative importance of the stock market in four of these eight cases, i.e., the US (24 lags, 75% in-sample) and Japan (12 and 24 lags, 50% in-sample; 12 lags, 75% in-sample). The four other monthly ARIMAs (two in the US and two in Germany) indicate a potential for lagged stock market information to emerge as a significant input into monthly IP growth forecasts, but the stock markets in both countries fail the AGS test. For the other four countries, the UK, Canada, France, and Italy, there are sixteen adequately estimated monthly models according to Table 3 and only two instances in Table 4 (UK, 24 lags, 50% and 75% in-sample) in which lagged stock returns lead to improved IP growth forecasts. At the monthly frequency, the AGS procedure imposes a stringent requirement on the prescience of the stock market. If past values of IP growth adequately predicts its future evolution, the stock market generally assumes a passive forecasting role. [8]

Regarding the prescience of the domestic stock market for domestic real activity in the G-7 countries, we offer possible explanations for our findings,

[8] The quarterly ARIMA results in Table 3 are not especially useful in interpreting the quarterly AGS results in Table 4 because the standard lag lengths for evaluating white noise residuals in ARIMAs are 12 and 24, which in the quarterly models represent 3 and 6 years, respectively. These lag lengths do not overlap the lengths of 4, 6 and 8 quarters for real stock returns that are used in the VARs in constructing the AGS tests.

which should be viewed as preliminary and speculative. The essence of the AGS procedure is the creation of IP growth forecasts from its own past. On a monthly and quarterly basis, the ARIMA results in Table 3 indicate that this task is relatively less problematic in France and Italy compared to the US and Japan. Over much of the out-of-sample forecasting periods considered, the French government was socialistic and nationalized large firms in several industries, while Italy operated with a large fraction of its industrial base under government control. Government-centered decision making tends to be less innovative and oriented towards stable rather than variable growth. In such an environment IP growth forecasts are likely to be highly correlated with past IP growth. Moreover, with a relatively large fraction of IP controlled by the government, the domestic stock markets in Italy and France would tend to be decoupled from the real sector. This latter feature may be especially relevant in interpreting the error-correction results for these two countries in Table 2. On the other hand, in innovative countries, such as the US, Japan, Canada and Germany, forecasting IP growth solely on the basis of its own past should be inherently more difficult, as is borne out in the monthly ARIMAs for Japan, Germany, and the US. The stock markets in these countries can play a more important role in gathering and evaluating information about future changes in IP growth that are not predictable from observations of its own past. This attribute seems to be especially borne out for Japan at the monthly frequency and for the US and Canada at the quarterly frequency. Germany is more problematic because it is uncertain whether a rational forecaster would have discovered the correct in-sample lag structure for stock returns. Finally, the behavior of the UK is also puzzling. At a monthly frequency, the ARIMA in Table 2 appears well-behaved, but the stock market is nonetheless informative about IP growth. Perhaps, the change to a conservative government and the slow dismantling of state-run enterprises in the UK over the out-of-sample period provide an opportunity for enhanced stock market forecasts of real activity. Other hypotheses, such as the influence of multinational firms, are topics for future research.

5. Summary and conclusion

This paper has examined the relationship between lagged real stock returns and the growth rate of industrial production for the G-7 countries using both in-sample cointegration and causality tests and an out-of-sample forecasting procedure introduced by Ashley et al. (1980). Our investigation was motivated by the Fama–Schwert results for the US in which in-sample goodness-of-fit tests are used to establish a relationship between industrial production growth and lagged real stock returns.

The results of our in-sample tests show that the log levels of industrial production and real stock prices are cointegrated in all G-7 countries. In addition, over a short-run horizon the error correction models indicate that the growth rate of industrial production is correlated with lagged real stock returns at some data frequency in six G-7 countries with Italy the exception. When the hypothesis of *improved* predictions of industrial production growth is analyzed with the out-of-sample AGS procedure, only the monthly results in the UK, Japan and perhaps the US and the quarterly results in Canada, the US and perhaps Germany show significant support for the hypothesis. Nonetheless, these findings tell us something that was not known previously. Prior research could not adequately address the issue of stock market prescience for the real sector because the estimation methods used are based on in-sample correlation, which could arise whether the stock market is prescient or not. The AGS procedure surmounts this stumbling block by comparing the out-of-sample forecasts of two competing algorithms. Stock market prescience can be identified with a significant relative improvement in forecast accuracy. We have found that the stock market is not prescient in every G-7 country because IP growth is sometimes so predictable that the stock market can make only a relatively minor contribution to understanding its future evolution. But we have also found evidence in the US, Canada, Japan, and the UK that domestic stock markets do incorporate information about future industrial production growth that is unavailable in only backward-looking data. In that sense, these G-7 stock markets are prescient for the real sector of their respective economies.

Acknowledgements

We wish to thank the referees for their very helpful comments and suggestions. An earlier version of this paper was presented at the 1997 Financial Management Association Meetings.

References

Ashley, R., Granger, C.W.J., Schmalensee, R., 1980. Advertising and aggregate consumption: An analysis of causality. Econometrica 48, 1149–1167.

Barro, R.J., 1990. The stock market and investment. Review of Financial Studies 3, 115–131.

Bittlingmayer, G., 1992. Stock returns, real activity, and the trust question. Journal of Finance 47, 1701–1730.

Box, G.E.P., Jenkins, P.M., 1970. Time Series Analysis, Holden Day, San Francisco.

Canova, F., Nicolo, G., 1995. Stock returns and real activity: A structural approach. European Economic Review 39, 981–1015.

Dickey, D.A., Fuller, W.A., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. Econometrica 49, 1057–1072.

Fama, E.F., 1990. Stock returns, expected returns, and real activity. Journal of Finance 45, 1089–1108.

Gallinger, G.W., 1994. Causality tests of the real stock return–real activity hypothesis. Journal of Financial Research 17, 271–288.

Geske, R., Roll, R., 1983. The fiscal and monetary linkage between stock returns and inflation. Journal of Finance 28, 7–33.

Granger, C.W.J., 1980. Testing for causality: A personal viewpoint. Journal of Economic Dynamics and Control 2, 329–352.

International Monetary Fund, International Financial Statistics, Washington, DC.

Kaul, G., 1987. Stock returns and inflation: The role of the monetary sector. Journal of Financial Economics 18, 253–276.

Lee, B.S., 1992. Causal relations among stock returns, interest rates, real activity, and inflation. Journal of Finance 47, 1591–1603.

Loflund, A., Nummelin, K., 1997. On stocks, bonds and business conditions. Applied Financial Economics 7, 137–146.

Park, S., 1997. Rationality of negative stock-price responses to strong economic activity. Financial Analysts Journal 53, 52–56.

Schwert, G.W., 1990. Stock returns and real economic activity: A century of evidence. Journal of Finance 45, 1237–1257.