UCLA
Recent Work

Title
Common Stock Returns and the Business Cycle

Permalink
https://escholarship.org/uc/item/9xz8m81x

Authors
Daniel, Kent
Torous, Walter

Publication Date
1991

Common Stock Returns and the Business Cycle[†]

Kent Daniel
Graduate School of Business
University of Chicago
1101 E. 58th Street
Chicago, IL 60637
312-702-4359

Walter Torous
Anderson Graduate School of Management
University of California, Los Angeles
Los Angeles, CA, 90024
310-825-2508

Abel's (1988) intertemporal asset pricing model implies that the autocorrelation pattern in expected returns reflects that observed in output growth rates. Consequently, by using the observed autocorrelation properties of macroeconomic data, we are able to provide univariate tests with power to detect deviations from the stationary random walk model over the post-World War II sample period. After regressing excess returns against industrial production's cyclical component, these univariate tests provide little evidence of serial correlation in the resultant residuals, confirming the presence of a business cycle effect in excess returns. However, our multivariate analysis concludes that while the business cycle contributes to these deviations from the stationary random walk model, predictable long term swings in expected returns arising from variable trends in macroeconomic data still remain.

Common Stock Returns and the Business Cycle

ABSTRACT

Abel's (1988) intertemporal asset pricing model implies that the autocorrelation pattern in expected returns reflects that observed in output growth rates. Consequently, by using the observed autocorrelation properties of macroeconomic data, we are able to provide univariate tests with power to detect deviations from the stationary random walk model over the post-World War II sample period. After regressing excess returns against industrial production's cyclical component, these univariate tests provide little evidence of serial correlation in the resultant residuals, confirming the presence of a business cycle effect in excess returns. However, our multivariate analysis concludes that while the business cycle contributes to these deviations from the stationary random walk model, predictable long term swings in expected returns arising from variable trends in macroeconomic data still remain.

I   Introduction

There is by now considerable evidence that stock returns are predictable. For example, Fama and French (1988) and Poterba and Summers (1988) provide empirical evidence of predictability in three- to five-year returns over the 1926-1987 sample period. Kim, Nelson, and Startz (1988), however, demonstrate that neither the Fama and French nor the Poterba and Summers test procedures reject the stationary random walk model over the post-World War II sample period. Similarly, Jegadeesh (1990) puts forward a more powerful test against the alternative hypothesis of an $AR(1)$ fads model of stock prices (Summers (1986)), but still fails to reject the stationary random walk hypothesis in post-World War II data.

While these univariate tests fail to reject the stationary random walk model, multivariate tests, in contrast, provide strong evidence against this null hypothesis over the post-World War II sample period. For example, Keim and Stambaugh (1986) and Fama and French (1989), among others, show that the default and term spreads reliably forecast stock and long-term bond returns.

That stock returns are predictable is not in itself evidence of market inefficiency. Expected returns should vary over time as intertemporal marginal rates of substitution, risk characteristics, and investors' risk tolerances change with prevailing economic conditions. Abel's (1988)

1

intertemporal asset pricing model, for example, relates excess returns to stochastically varying output growth. Consistent with this, Chen (1991) presents evidence suggesting that the ability of the default and term spreads to predict excess returns is due to their being business cycle indicators that provide information about current and future economic conditions.

In light of this multivariate evidence, it is surprising that univariate tests remain inconclusive by not taking advantage of the economic explanations offered for multivariate predictability. Of course, even if expected returns do vary with prevailing economic conditions, excess returns need not exhibit univariate predictability if changes in economic conditions themselves are serially uncorrelated. However, the observed characteristics of such economic aggregates as GDP and unemployment suggest that this is not the case. As noted by Zarnowitz (1992, p. *xv*), "business cycles are characteristically persistent and pervasive." Appealing to Abel's model, expected returns should exhibit this same persistence and, as a result, we can exploit the observed autocorrelation properties of macroeconomic data to design univariate tests which reliably reject the stationary random walk model over the post-World War II sample period. By contrast, we demonstrate that previous univariate tests lack power to detect deviations brought about by business cycle related variation in expected returns.

Abel's model further implies that any univariate predictability in excess returns should be explained by corresponding univariate predictability in output growth and its volatility. Thus, the component of current and past output which predicts future growth should also predict future excess returns. However, predicting future output growth from current and past data is complicated by the fact that output, proxied in this paper by industrial production, is an integrated process ($I(1)$). Econometrically, we overcome this problem by decomposing this time series into its permanent and temporary components. Beveridge and Nelson (1981) establish that any integrated series can be decomposed into a pure random walk with drift, interpreted as its trend component, and a stationary component, interpreted as its cyclical component. Consistent with Lucas (1977), this specification views the business cycle as stationary movements about a stochastic trend.[1] To the extent that production's cyclical component provides a good proxy for future production growth, Abel's model suggests that, after controlling for their cyclical component, excess returns should be serially uncorrelated.

2

However, errors can arise if this cyclical component is estimated by the residuals of a fitted linear trend (Nelson and Kang (1981) and Nelson and Plosser (1982)). First differencing a time series with both growth (non-stationary) and cyclical (stationary) components is also inappropriate as this simply confounds the contribution of each component (Stock and Watson (1988)). We avoid these difficulties by using Watson's (1986) unobserved components (*UC-ARIMA*) method as well as a one-sided version of the Hodrick and Prescott (1980) filter to stochastically detrend post-World War II industrial production data. The robustness of our results are also verified using spectral band-pass filters (Baxter and King (1994)).

We confirm, using long-horizon regressions, that industrial production's cyclical component is inversely related to excess returns. Furthermore, after regressing excess returns against production's estimated cyclical component, we find that, consistent with Abel's model, our univariate test procedures provide little evidence against the null hypothesis that the resultant residuals are serially uncorrelated. This result complements Ferson and Harvey's (1991) conclusion that most of the time variation in excess returns may be attributed to a single factor. However, we find that these residuals are themselves predictable using the default and term spreads. According to Abel's model, these non-production variables can forecast excess returns *only* if they provide information about future production growth or production volatility. Consistent with this, we provide evidence that the term spread is primarily informative about future growth prospects while additional information about future production volatility is contained in the default spread.

The plan of this paper is as follows. We begin by discussing Abel's intertemporal asset pricing model and highlight how this framework motivates our subsequent empirical analyses. In Section III we develop our univariate test procedures. Asymptotically and in small samples, these tests have significantly more statistical power than previous tests in detecting business cycle deviations from the stationary random walk model. In Section IV we characterize the comovement of excess returns with the cyclical component of industrial production as well as its stochastic trend. We demonstrate that the rejection of the stationary random walk hypothesis in post-World War II data is, to a large extent, attributable to countercyclical movements in excess returns. We also use our time series decomposition of industrial production into its

3

trend and cycle components to further investigate the relation between excess returns and business cycle indicators, such as the term and default spreads. Section V provides a summary and conclusions.

II Expected Return Variation Over the Business Cycle

Abel (1988) provides a multiperiod general equilibrium asset pricing model in which dividend growth and volatility evolve stochastically over time. For a log utility representative individual (see also Chen (1991)), Abel shows that the risk free rate and the expected market risk premium are given by

$$R^f_{t+1} = \beta^{-1} \left(\frac{\mu_t}{y_t} \right) \cdot \frac{1}{1 + \nu_t^2} \tag{1}$$

and

$$E_t\{R^e_t\} = E_t\{R_{t+1} - R^f_{t+1}\} = \beta^{-1} \left(\frac{\mu_t}{y_t} \right) \frac{\nu_t^2}{1 + \nu_t^2} \tag{2}$$

where

$$
\begin{aligned}
\beta &\equiv \text{ time invariant discount factor} \\
y_t &\equiv \text{ the current level of output} \\
\mu_t &\equiv E_t[y_{t+1}] \equiv \text{ the expected level of future output} \\
\nu_t &\equiv var_t(y_{t+1})/\mu_t^2 \equiv \text{ the conditional coefficient of} \\
&\qquad \text{variation of future output.}
\end{aligned}
$$

Notice from (1) and (2) that the risk free rate and the expected market risk premium are countercyclical. That is, both are decreasing in the current level of production y_t for a *given* level of future production, μ.[2] Intuitively, in an expansion (recession) when income is temporarily high (low) relative to wealth, the representative individual is more (less) willing to substitute current consumption for future consumption and, all else equal, returns are correspondingly low (high). According to (2), the expected risk premium is also increasing in the volatility of production, ν^2.

Abel's model implies that excess returns are predictable only to the extent that future production growth and production volatility are themselves predictable. For example, if we assume

4

that production volatility is constant, then persistence in excess returns requires persistence in production growth.

To further understand the implications of Abel's model for the time series behavior of excess returns, we follow, among others, Beveridge and Nelson (1981) and Watson (1986) and decompose the log of production into a permanent or *stochastic trend* component (τ_t) and a temporary or *cyclical* component (c_t)[3]:

$$\ln y_t = \tilde{\tau}_t + \tilde{c}_t. \tag{3}$$

The trend component $\tilde{\tau}_t$ is modeled as a random walk with drift,

$$\tilde{\tau}_t = \delta + \tilde{\tau}_{t-1} + \tilde{e}_t^\tau, \qquad var(\tilde{e}_t^\tau) = \sigma_\tau^2, \tag{4}$$

while the temporary or cyclical component, $\tilde{c}_t$ is characterized by a stationary ARMA process

$$\phi(L)\tilde{c}_t = \theta(L)\tilde{e}_t^c, \qquad var(\tilde{e}_t^c) = \sigma_c^2, \tag{5}$$

where $\phi(L)$ and $\theta(L)$ denote finite order polynomials in the lag operator L, with stationarity requiring that $\theta(L)/\phi(L)$ have roots outside the unit circle.[4]

Innovations in production's trend, $\tilde{e}_t^\tau$, are, by construction, unforecastable using current and past production data, while future changes in the stationary component of production, Δc_{t+j}, are forecastable using current and past production data. Innovations in both the trend and cycle ($\tilde{e}_t^\tau$ and $\tilde{e}_t^c$), however, may be forecastable if we expand our information set to include non-production data, for example, the default spread or the term spread. Abel's model implies that these variables can forecast excess returns *only* if they provide information about future production growth or production volatility.

If production is a pure trend (*i.e.*, a random walk), then future production growth will be uncorrelated with current and past production growth, and, from (2), the expected (log of the) market risk premium will be unrelated to past and current economic conditions. However, if, as is more likely, production is characterized by a persistent cyclical component, then expected future growth will, on average, be high (low) when this temporary component is relatively low (high), implying that the expected market risk premium will be correspondingly high (low).

Abel's model also has implications for the serial correlation properties of excess returns. Again, if production follows a random walk, then, since production growth is uncorrelated,

5

excess returns will also be serially uncorrelated. Alternatively, the presence of a persistent cyclical component in production implies that past and future production growth will be correlated, and, as a result, excess returns will now be serially correlated. In addition, Abel's model implies that the pattern of serial correlation in excess returns should be closely related to the pattern of serial correlation in production growth.

We make use of these implications of Abel's model in several ways. First, since extant empirical evidence confirms the presence of a cyclical component in production, Abel's model implies that excess returns should also be serially correlated. Furthermore, since the autocorrelation pattern in expected returns should be closely related to that observed in production growth rates, this allows us to design univariate tests which have power to detect deviations from a stationary random walk model for stock prices based upon the observed autocorrelation properties of production data. Additionally, if we can identify production's cyclical component then it should predict future excess returns. To the extent that production's cyclical component is a good proxy for future production growth, the component of returns that remains once we control for its cyclical component should be serially uncorrelated, although it may very well be predictable using other variables in the available information set.

III Univariate Tests

From Abel's model, risk premia vary counter-cyclically with the business cycle, implying that the time series behavior of production growth rates should be reflected in the stochastic properties of excess returns. In this section, we develop univariate test procedures which have statistical power, both asymptotically and in small samples, against this alternative hypothesis.

To develop such tests requires that we exploit the known properties of the business cycle. While many alternative definitions of the business cycle are available (for example, see Sargent (1987), especially pages 279-283), an exact characterization is, unfortunately, difficult. Furthermore, misspecifying the alternative may result in a less powerful test.

One business cycle restriction which may reasonably be imposed is that movements in economic aggregates are persistent, or more formally, that their spectral density at high fre-

6

quencies is low. Intuitively, persistence implies that if we are in an expansion (or recession) today, we are likely to still be in an expansion (or recession) six months from today. For example, Zarnowitz (1992, p. 22) writes

> The observed fluctuations vary greatly in amplitude and scope as well as duration, yet they also have much in common. First, they are national, often international, in scope, showing up in a multitude of processes, not just in total output, employment and unemployment. Second, they are persistent - lasting, as a rule several years, that is, long enough to permit the development of cumulative movements in the downward as well as upward direction.

The persistence of the business cycle motivates our univariate test procedures. Abel's model implies that returns and, as we demonstrate in Appendix A, return autocorrelations should exhibit this same persistence. However, this persistence will typically be obscured by noise in the underlying data. From spectral techniques, we diminish the effects of this noise by applying a moving average filter to return autocorrelations. By doing so, we reduce this noise without significantly altering the business cycle induced pattern of persistent autocorrelations.[5]

A   The Power of the Fama and French Regression Test

We first provide intuition why previous tests lack power to detect persistence in stock return autocorrelations. For expositional purposes, we couch our analysis in the context of Fama and French's (1988) regression test.[6]

The Fama and French test examines whether $\beta(\tau) = 0$ in the following regression:

$$R(t, t + \tau) = \alpha(\tau) + \beta(\tau) \cdot R(t - \tau, t) + \epsilon(t, t + \tau)$$

where $R(t, t + \tau)$ represents the stock's return from t to $t + \tau$. Their ordinary least squares slope estimator is:

$$\hat{\beta} = \frac{\frac{1}{T} \sum_t r(t, t + \tau) r(t - \tau, t)}{\frac{1}{T} \sum_t r^2(t - \tau, t)},$$

where r denotes corresponding demeaned returns. The asymptotic power of the Fama and French regression test may be analytically derived for a local alternative hypothesis using the arguments of Davidson and MacKinnon (1987).[7] Recall that for a sufficiently large sample size, $(T \to \infty)$, we can reject a null hypothesis with certainty for any fixed alternative hypothesis. Therefore, to properly evaluate the asymptotic power of the Fama and French test, we consider

7

the limit of a Pitman sequence of return series for which the distance between the alternative and null hypotheses becomes infinitesimal as the sample size goes to infinity. The rate at which the alternative converges towards the null is such that in the limit the probability of rejecting the null, when the alternative is true, lies in the interval $(0,1)$ as $T \to \infty$. This requires that the variance of the temporary component of common stock prices goes to zero as $T \to \infty$. It is clear that for the limit of this sequence, the probability limit of the denominator will be:

$$\frac{1}{T}\sum_t r^2(t-\tau, t) = \frac{1}{T}\sum_t \sum_{s=-\tau}^{s=\tau} |s-\tau| r_t r_{t-s} = \frac{\tau}{T}\sum_t r_t^2.$$

Similarly, the probability limit of the numerator will be:

$$\frac{1}{T}\sum_t r(t, t+\tau) r(t-\tau, t) = \frac{1}{T}\sum_t \left(\sum_{u=t}^{t+\tau-1} r_u\right)\left(\sum_{u=t-\tau}^{t-1} r_u\right)$$

$$= \frac{1}{T}\sum_t \sum_{s=1}^{2\tau-1} min(s, 2\tau-s) r_t r_{t+s}$$

$$= \sum_{s=1}^{2\tau-1} min(s, 2\tau-s) \frac{1}{T}\sum_t r_t r_{t+s}.$$

Therefore, the *plim* of the slope estimator will be

$$\hat{\beta} = \frac{1}{\tau}\sum_{s=1}^{2\tau-1} min(s, 2\tau-s), \hat{\rho}_s \tag{6}$$

showing that the Fama and French regression test at return horizon τ is asymptotically equivalent to a test of the hypothesis that a weighted sum of return autocorrelations between lags 1 and $2\tau-1$ is equal to zero.

Notice that Fama and French's regression test implicitly assumes a triangular weighing function of return autocorrelations. Whether this provides a powerful test of the random walk hypothesis depends upon the pattern of return autocorrelations expected under the alternative hypothesis. We show in Section B that under the null and local alternative hypotheses, autocorrelation estimators at different lags are uncorrelated and have approximately the same standard error. Given this, the most powerful test statistic will weight sample autocorrelations so that the weight at a particular lag length will be proportional to the corresponding autocorrelation coefficient expected under the alternative hypothesis. As a result, the Fama and French regression test is likely to have little power against reasonable alternative hypotheses positing slow variation in expected returns over the business cycle.

Since a test's power depends critically on the nature of the alternative hypothesis, a powerful test must impose *a priori* reasonable restrictions on the behavior of expected returns over the

8

business cycle. A general test, such as Box and Pierce's (1970) portmanteau or Q test, will lack power precisely because it imposes little or no such restrictions.[8] On the other hand, it may also be inappropriate to rely on a test positing an exact parameterization of returns over the business cycle, for example, an $AR(p)$ specification. While such a test will more likely reject the stationary random walk model if such a specification actually characterizes the expected return generating process, it may have little or no power if expected returns follow some other equally reasonable alternative.[9]

B Powerful Autocorrelation Based Tests

Consider a demeaned stationary time series $\{r_t\}_{t=1}^T$ and define its autocorrelogram by

$$\rho_\tau = \frac{Cov(r_t, r_{t+\tau})}{Var(r_t)} \qquad \text{for } \tau > 0. \tag{7}$$

The autocorrelogram summarizes the autocorrelations of a stationary time series as a function of lag length τ.

We estimate the autocorrelation coefficient at lag length τ by

$$\hat{\rho}_\tau = \frac{T \sum_{j=1}^{T-\tau} r_j r_{j+\tau}}{(T-\tau) \sum_{j=1}^T r_j^2}. \tag{8}$$

However, to minimize the effects of any noise in the data, we consider the averaged autocorrelogram estimated by

$$\hat{\rho}_\tau^* = \frac{1}{2\delta + 1} \sum_{i=\tau-\delta}^{\tau+\delta} \hat{\rho}_i. \tag{9}$$

That is, the averaged autocorrelation at τ is estimated by the arithmetic average of the corresponding $2\delta + 1$ sample autocorrelation coefficients centered at lag length τ. Under the null hypothesis that $\{r_t\}_{t=1}^T$ is serially uncorrelated, it follows that

$$\sqrt{(2\delta + 1) \cdot (T-\tau)} \cdot \hat{\rho}_\tau^* \sim \mathcal{N}(0, 1). \tag{10}$$

The motivation for this averaging is similar to that of smoothing or windowing in spectral analysis.[10] However, if δ is chosen too large, the resultant over-averaging will eliminate any long-term persistence. Alternatively, a too small δ value will not minimize the effects of any noise. We use $\delta = 3, 4$, and 5. Our choice follows directly from the properties of the business

cycle tabulated by macroeconomists. In particular, peacetime expansions (from trough to peak) in the United States over the period 1933 to 1982 have an average duration of 37 months with a standard deviation of 15 months, but peacetime contractions (from peak to trough) have only an average duration of 11 months with a standard deviation of 3 months (see Zarnowitz (1992) p. 23, Table 2.1). As a result, by choosing $\delta = 3, 4$, and 5, the averaging window's width is approximately 11 months, ensuring that we do not average over more than one contraction.

To assess whether sampled averaged autocorrelations provide statistical evidence against the stationary random walk hypothesis, we provide two univariate test procedures to detect business cycle induced deviations. The first, the *Sparse* χ^2 or χ^2_S test statistic,

$$\chi^2_S = (2\delta + 1) \sum_{j=1}^{N} (T - (\tau_j + \delta)) \hat{\rho}^{*2}_{\tau_j} \tag{11}$$

where $\tau_j = \delta + (j - 1)(2\delta + 1)$, directly tests the joint restriction imposed by the stationary random walk model across the *averaged* autocorrelogram

$$H_0 : \qquad \rho^*_{\tau_1} = \ldots = \rho^*_{\tau_N} = 0.$$

The derivation of the χ^2_S test statistic and its small sample properties are detailed in Appendix B.1.

A potential shortcoming of the χ^2_S statistic is that its realized value may be sensitive to the particular averaging interval chosen. In other words, since the averaged autocorrelogram is infrequently (or sparsely) sampled via expression (11), significant departures from the stationary random walk model evidenced, for example, by a peak in the averaged autocorrelogram, may not be detected.

Consequently, we also rely on a second univariate test procedure, the Bonferroni test statistic, which is not subject to this potential shortcoming. In particular, consider the set of averaged autocorrelations coefficients $\{\hat{\rho}^*_\tau\}$ generated by expression (9).[11] Using the Bonferroni inequality, we can calculate an upper bound for the probability, under the null hypothesis H_0, of an averaged autocorrelation being at least as extreme as the most extreme of the observed averaged autocorrelations, $\max_\tau |\hat{\rho}^*_\tau|$. This upper bound is independent of the correlation structure of $\{\hat{\rho}^*_\tau\}$. If this probability is less than conventional significance levels, we may confidently reject H_0.

10

The derivation of the Bonferroni test statistic and its small sample properties are detailed in Appendix B.2. Appendix B.3 investigates the empirical power of both of our univariate tests and demonstrates their power in small samples to detect business cycle deviations from the stationary random walk model.

C Univariate Test Results

We apply our test procedures to monthly excess returns over the post-World War II sample period, 1947:1 - 1992:12. The results are tabulated in Table I.

Notice that in the absence of averaging, $\delta = 0$, both the χ_S^2 (Panel A) and Bonferroni tests (Panel B) reject the random walk model for the small firm portfolio, decile 1. Given the relative importance of these small stocks to the equal-weighted (EW) CRSP portfolio, we see that for $\delta = 0$ the χ_S^2 statistic also rejects the random walk model for EW excess returns.[12] However, without averaging, we cannot reject the random walk model for excess returns to either the value-weighted (VW) CRSP portfolio or the large firm portfolio, decile 10.

With averaging, both the χ_S^2 and Bonferroni tests now provide evidence against excess returns of the EW and VW portfolios following a random walk. But as suggested earlier, the performance of the χ_S^2 test appears sensitive to the choice of δ and the width of the resultant averaging window. In particular, while the Bonferroni test rejects EW and VW excess returns following a random walk for each of $\delta = 3, 4$, and 5, the χ_S^2 test only rejects for $\delta = 4$. In addition, both tests reject the null hypothesis with greater significance for the VW portfolio. This result may reflect the greater volatility of excess returns to small firm portfolios. Similar conclusions hold for the decile portfolios where, in particular, the Bonferroni test allows us to reject the random walk hypothesis for the large firm portfolio.[13]

IV Multivariate Tests

Our preceding empirical analysis demonstrates that excess returns do not follow a stationary random walk. Since our univariate test procedures, motivated by Abel's model, are designed to exploit the persistence which characterizes the business cycle, this suggests the presence of a business cycle effect in excess returns. To statistically confirm and characterize this effect,

11

we now provide a multivariate empirical analysis of excess returns and the business cycle.

Our multivariate analysis is related to that of Chen (1991) who investigates whether the ability of the default and term spreads, as well as other variables including the market dividend yield, to forecast excess returns is due to their power to forecast future output growth and output volatility. However, we extend Chen's analysis in several ways.

First, Chen confirms that these forecasting variables are indicators of current and future economic prospects and can indeed forecast excess returns. Yet it may very well be the case that a different component of these variables is forecasting excess returns than is forecasting output growth or volatility. Therefore, it is important to investigate whether there exists a *direct* link between output and excess returns. To do so, however, is complicated by the fact that measures of output, such as industrial production, are integrated variables, requiring use of stochastic detrending methods to properly obtain their stationary component.

Secondly, motivated by Fama and French's (1989) finding that the default spread and the market dividend yield appear to track longer-run variability in economic conditions than does the term spread, we also investigate the temporal characteristics of these forecasting relationships. We do so by using spectral techniques, as discussed by Baxter and King (1994), to decompose excess returns, the forecasting variables, and production growth and its volatility into their respective trend, cycle, and high frequency components.

The sample period used in our multivariate tests begins in 1953:01. Industrial production data is available on a monthly basis from CITIBASE beginning in 1947:01, but six years of monthly data (72 observations) are used to initialize our stochastic detrending methods.

A Stochastic Detrending of Industrial Production

We use two stochastic detrending methods to empirically estimate the temporary component of post-Korean War industrial production: Watson's (1986) Unobserved Components ARIMA (*UC-ARIMA*) method and the Hodrick and Prescott (1980) filter.

12

A.1 The *UC-ARIMA* Method

The *UC-ARIMA* method recognizes that while the log of industrial production, lny_t, is itself observable, its additive components, τ_t and c_t, are individually unobservable. More significantly, to identify parameters, the *UC-ARIMA* method decomposes the observed series assuming that the trend and stationary innovations are uncorrelated:

$$cov(e_t^\tau, e_{t-k}^c) = 0 \qquad \forall\, k.$$

That is, the economic factors giving rise to trend innovations are assumed to be unrelated to the economic sources of business cycle movements.

Maximum likelihood estimation of the *UC-ARIMA* model is carried out by casting the model into state-space form and using the Kalman filter initialized at a vague prior (Harvey (1981)). Given the maximum likelihood parameter estimates, the Kalman filter also yields corresponding estimates of the unobservable state variables (τ_t, c_t). The maximum likelihood estimation results for the 1953:1 to 1992:12 sample period are as follows (with asymptotic standard errors in parentheses)

$$
\begin{aligned}
\ln y_t &= \tau_t + c_t \\
\Delta\tau_t &= \underset{(.00031)}{.00315} + e_t^\tau & \sigma_\tau &= \underset{(.00058)}{.00608} \\
c_t &= \underset{(0.014)}{1.664} \cdot c_{t-1} - \underset{(0.014)}{0.6848} \cdot c_{t-2} + e_t^c & \sigma_c &= \underset{(0.00060)}{0.00652}.
\end{aligned}
$$

These results are consistent with an average monthly growth rate of 0.32% in industrial production over our sample period. Industrial production's temporary component is estimated by an AR(2) specification; higher order autoregressive coefficients were found to be statistically insignificant.[14]

A.2 The Hodrick-Prescott Filter

To ensure that our subsequent empirical results are not specific to a particular stochastic detrending method, we also apply the Hodrick-Prescott filter to our industrial production data. The Hodrick-Prescott filter uses spline smoothing (Reinsch (1967)) to stochastically detrend macroeconomic data. We modify their procedure to utilize only past data (through time t) to estimate industrial production's stochastic trend at time t. This one-sided version of

13

the Hodrick-Prescott filter ensures that any empirically observed relationship between realized returns and the estimated cyclical component does not simply reflect the empirical fact that stock returns forecast future changes in industrial production (Fama (1981, 1990)).

Our one-sided version of the Hodrick-Prescott filter uses the previous $n = 120$ monthly observations of industrial production to estimate τ_t. For each t we determine the function $\{\hat{\tau}_s\}_{s=t-n}^{t}$ which for a given λ minimizes

$$\sum_{s=t-n}^{t} (\ln y_s - \hat{\tau}_s)^2 + \lambda \sum_{s=t-n+1}^{t-1} [(\hat{\tau}_{s+1} - \hat{\tau}_s) - (\hat{\tau}_s - \hat{\tau}_{s-1})]^2 . \tag{12}$$

The trend at time t is then estimated by $\hat{\tau}_t$ and the estimate of the cyclical component is given by $\hat{c}_t = \ln y_t - \hat{\tau}_t$.

Here λ represents the trade-off between closeness of fit, measured by the residual sum of squares (the first portion of expression (12)), and the smoothness of τ_t, measured by the integrated squared second derivative of the trend (the second portion of expression (12)).

As noted by King and Rebello (1989), this stochastic detrending method can also be interpreted as a moving average filter which for appropriate λ values removes low frequency movements in the underlying data. We chose that value of λ (3×10^6) which removes fluctuations having periodicity of greater than 8 years in monthly industrial production data. Recall that according to NBER calculations, the business cycle in post-World War II macroeconomic data has periodicity of between approximately two and eight years. This also corresponds to the stochastic characteristics of the Hodrick-Prescott filter implemented by Kydland and Prescott (1990), as well as by Hodrick and Prescott (1980) themselves.

A.3   Industrial Production's Estimated Temporary Component

In Figure 1 we plot the estimated temporary component obtained by applying our one-sided Hodrick-Prescott filter to industrial production data over the 1953:01 to 1992:12 sample period. For comparison purposes, we also plot the corresponding temporary component estimated by the *UC-ARIMA* method as well as the NBER datings of business cycle peaks and troughs.[15] Notice that regardless of the stochastic detrending method used, there appears to be broad correspondence between our estimates and the NBER's peak to trough business cycle periods.

14

B Excess Returns and Industrial Production's Temporary Component

According to Abel's model and our time series specification of industrial production dynamics, we should observe an inverse relationship between excess returns and industrial production's temporary component. Furthermore, if the previously documented deviations from the stationary random walk model are due to business cycle related movements in expected returns, then once we control for industrial production's cyclical component, our univariate tests should provide little, if any, evidence against the null hypothesis that the resultant residuals are serially uncorrelated. Table II presents the results of regressing excess returns to the EW, VW, and size decile portfolios against industrial production's temporary component estimated by the UC-ARIMA method (Panel A) and the Hodrick-Prescott filter (Panel B). Consistent with countercyclical movements in expected returns, notice that the regression slope coefficients are significantly negative throughout. The smaller firm portfolios appear to be more sensitive to business cycle movements than the larger firm portfolios. However, the explanatory power of these regressions, as measured by their adjusted $R^2 s$, is low. Assuming that all of the business cycle variation in expected returns has been captured, these results suggest that this business cycle variation is small when compared to the total variation observed in realized monthly returns. To determine whether the previously documented deviations from the stationary random walk model are attributable to business cycle related movements in expected returns, we investigate the stochastic behavior of these excess returns once the estimated cyclical effect is controlled for. We do so by subjecting the residuals of Table II's regressions to our univariate test procedures. The results are presented in Table III where it is evident that the null hypothesis of serially uncorrelated residuals is rejected less frequently than the corresponding null hypothesis for excess returns. In fact, the tests provide no evidence of persistence in the EW and VW residuals at conventional significance levels.[16] This result compliments Ferson and Harvey's (1991) conclusion that most of the time variation in returns may be attributed to a single factor.

Unfortunately, this analysis ignores the fact that stochastic detrending methods estimate industrial production's temporary component with error. Consequently, an errors-in-variable problem may arise. To ensure the robustness of the results presented in Table III, we generalize

our UC-ARIMA framework as follows:

$$
\begin{aligned}
r_t &= \alpha + \beta c_t + \epsilon_t^r \\
\ln(y_t) &= \tau_t + c_t \\
\tau_t &= \delta + \tau_{t-1} + \epsilon_t^\tau \\
c_t &= \rho_1 c_{t-1} + \rho_2 c_{t-1} + \epsilon_t^c
\end{aligned}
\tag{13}
$$

where, in addition, we assume that

$$
cov(e_t^r, e_{t-k}^c) = cov(e_t^r, e_{t-k}^\tau) = cov(e_t^c, e_{t-k}^\tau) = 0 \ \forall \ k.
$$

This joint estimation allows us to investigate the relationship between excess returns and industrial production's temporary component while explicitly recognizing that this cyclical component is estimated with error. Table IV presents the results of applying our univariate tests to the residuals for the EW, VW, and size decile 1 and 10 portfolios. As in Table III, we still see little evidence of persistence in these residuals.[17]

B.1 Business Cycle Indicators and the Predictability of Excess Returns

The preceding analysis provides evidence that the temporary component of industrial production contributes directly to the deviations from the stationary random walk model detected by our univariate test procedures. This conclusion is consistent with the evidence of Fama and French (1989) and others that the default and term spreads can reliably forecast excess returns, especially at long horizons. Default and term spreads are business cycle indicators, with both spreads being low around peaks and high around troughs. If expected returns also have a business cycle component, high returns when business conditions are poor and low returns when business conditions are good, then the default and term spreads are tracking this predictable variation in expected returns.

The estimated cyclical component of industrial production also reliably forecasts excess returns. This can be seen in Table V where long-horizon regressions are used to investigate the predictability of excess EW (Panel A) and excess VW (Panel B) returns using the UC-ARIMA estimated cyclical component as well as the default and term spreads.[18] Like the term spread, the estimated cyclical component only forecasts over approximately the next 12

months, with excess VW returns being more predictable than excess EW returns. This latter result is consistent with our univariate test procedures providing stronger evidence against the random walk hypothesis for excess VW returns. In contrast, the default spread forecasts excess returns over longer return horizons. The explanatory power of the default spread is particularly impressive for long horizon excess VW returns; for example, the default spread alone explains approximately 30% of the variation in 48 month excess VW returns.

If this evidence of predictability is consistent with the term spread and the default spread tracking the business cycle variation in excess returns, then these spreads should track industrial production's cyclical component itself. Table VI directly investigates this by using the UC-ARIMA estimated cyclical component and regressing its changes over successive past and future quarters against the default and term spreads. The countercyclical behavior of both these spreads is readily apparent. For example, the term spread is significantly related to both past and future changes in the estimated cyclical component of industrial production. As measured by the regressions' R^2s, the term spread is particularly informative about future cyclical changes (up to four quarters). In comparison, while the default spread is also significantly related to both past and future changes in the estimated cyclical component, it appears to provide only a modest amount of information about these changes (up to four quarters in either direction).

However, by construction, changes in production's cyclical component (Δc) do not reflect changes in its trend component ($\Delta \tau$) which also contribute to future production growth (Δlny). To capture this, we regress changes in the log of industrial production over successive past and future quarters against the term and default spreads as well as its estimated cyclical component (Table VII). To ensure the robustness of our results, we also use gross domestic product (GDP) (Table VIII).

Like Chen (1991), we find that the default spread is only correlated with immediate past growth rates in industrial production and GDP, although the informativeness of these regressions, as measured by their corresponding R^2s, is minimal. In contrast, the term spread is correlated, in a countercyclical fashion, with both future and, unlike Chen, past growth rates.[19] These regressions are particularly informative about the future growth of the economy (up to four quarters). Similar to the default spread, the estimated cyclical component is only corre-

lated with past growth rates in industrial production and GDP. However, unlike the default spread regressions, these regressions are quite informative about these past growth rates. That the cyclical component provides little information about the economy's future prospects is not surprising since it is estimated on the basis of current and past industrial production data. We also use long horizon regressions to investigate whether the residuals from regressing excess VW returns on the UC-ARIMA estimated cyclical component are themselves predictable using the default and term spreads as well as the estimated cyclical component (Table IX). Like the long horizon excess return regressions, we find that the term spread is significant in forecasting these residuals out to a 12 month horizon, while the default spread is significant out to a 48 month horizon. Also, the explanatory power of the default spread in these regressions significantly exceeds that of the term spread. The estimated cyclical component, as expected, is not found to be significant throughout.

While the default spread can reliably forecast excess returns (Table V), it says little about the current health of the economy or its future growth prospects (Table VI, VII, and VIII). Appealing to Abel's model, this suggests that the default spread provides information about production's volatility. Following Chen (1991), we proxy this volatility by the absolute value of the residuals obtained from regressing changes in the log of industrial production against the average of the default spread over the preceding three months as well as last month's term spread and cyclical estimate[20]:

$$\Delta \log(IP_t + j) = \alpha + \beta_1 \sum_{i=1}^{3} DEF_{t-i} + \beta_2 TERM_{t-1} + \beta_3 Cycle_{t-1} + \epsilon_t + j$$

and our proxy for volatility is[21]:

$$\hat{\sigma}_{IP} = |\epsilon_t|.$$

We then regress changes in production volatility over successive past and future quarters against the term spread and the default spread. The results are tabulated in Table X. Notice that the default spread and, to a lesser extent, the term spread are correlated with future changes in production volatility. In particular, a high (low) default spread today is associated with an increase (decrease) in future production volatility, while a high (low) term spread today implies a decrease (increase) in future production volatility.

Consistent with Chen (1991), we conclude that the term spread provides information about

18

the economy's future growth prospects. However, while Chen attributes the default spread's ability to forecast excess returns to its informativeness regarding the economy's current health, our results, in contrast, demonstrate that the default spread's forecasting ability arises primarily from its informativeness regarding future production volatility. We explore the macroeconomic determinants of this predictive ability in the next section.

B.2 Frequency Band Correlations

The first differencing of macroeconomic variables, such as production or its volatility, potentially obscures any long term dependence which may exist between these variables and excess returns by removing much of this data's low frequency component. Such long term dependence has previously been suggested by Fama and French (1989) who argue, using long horizon regressions and by visual inspection of the data, that long-run variation in expected returns is captured by the default spread, while the term spread forecasts expected returns' shorter-run variation.

In this section, we directly investigate this hypothesis by separately correlating movements in excess returns with movements in default and term spreads at frequencies corresponding to trend movements, business cycle movements, as well as the remaining high-frequency movements in the data. To do so, we filter these variables to isolate their (i) fluctuations which exceed 32 quarters in length (trend component); (ii) fluctuations between 6 and 32 quarters in length (business cycle component); and (iii) fluctuations less than 6 quarters in length (high-frequency component). As noted by Baxter and King (1994), this specification of frequency bands is consistent with Burns and Mitchell's (1946) finding that business cycles tend to be no less than 6 quarters in length but typically last fewer than 32 quarters.[22]

Table XI presents these frequency band correlations for excess EW and VW returns. Notice that excess returns are significantly positively correlated with the default and term spreads at the business cycle frequencies. In other words, consistent with the previous multivariate empirical evidence of Keim and Stambaugh (1986), Fama and French (1989) and others, these spread variables *do* capture the business cycle component of excess returns, with high (low) default and term spreads being significantly associated with high (low) excess returns at these frequencies. However, excess returns are correlated with the default and term spreads at

19

other frequencies. At high frequencies we see a significant positive correlation only between excess returns and the term spread. In contrast, at the trend frequencies, a significant positive correlation is found only between the default spread and excess returns. Therefore, consistent with Fama and French's claim, the default spread's trend or low frequency component reliably tracks excess returns' low frequency component, while high frequency movements in excess returns are captured by the term spread's high frequency component.

To better understand the macroeconomic determinants of their predictive ability, Table XI also gives the correlations of the filtered components of the term and default spreads with production growth rates and production volatility. Production growth rates are significantly positively correlated with the term and default spreads at the business cycle frequencies, with the term spread's correlation being slightly higher. However, these spreads are only marginally correlated with extremely short run variation in production growth rates, and appear to provide no information about the long run variation in production growth rates. Alternatively, we see that production volatility appears to be significantly and positively correlated with the default spread at all frequency bands. In particular, a positive correlation obtains between production volatility and the default spread at the trend frequencies.[23]

The preceding spectral analysis provides an interesting perspective on the relationship between stock returns and the macroeconomy. The dynamics of macroeconomic variables, such as production or its volatility, are characterized by both a stochastic trend component as well as a business cycle component. Appealing to Abel's model, excess returns will also be characterized by stochastic trend and business cycle components; it is this latter component which is detected by our univariate test procedures as business cycle deviations from the stationary random walk model. However, the default and term spreads provide additional information about excess returns beyond that contained at business cycle frequencies. In particular, long swings in excess returns, attributable to changes in the macroeconomy's trend, are captured by the default spread whose long run predictive ability arises from the fact that it contains information about corresponding long swings in production volatility.

20

V Conclusions

Multivariate tests of market efficiency, for example, Keim and Stambaugh (1986) and Fama and French (1989), conclude that business cycle indicators, such as the default spread and the term spread, reliably forecast stock returns. This result is generally attributed to the systematic variation in expected returns over the business cycle.

Since the distinguishing characteristic of the business cycle is its persistence, Abel's (1988) intertemporal asset pricing model implies that this same persistence will also be evident in the time series behavior of expected returns. Unfortunately, extant univariate tests of market efficiency provide little power against reasonable alternative hypotheses positing slow variation in expected returns over the business cycle. By contrast, this paper's univariate test procedures, motivated by Abel's model and based on a weighted autocorrelogram of returns, allow us to reject, with a high degree of statistical confidence, the null hypothesis that expected returns are constant over the business cycle.

We confirm that industrial production growth rates and excess returns exhibit similar autocorrelation patterns in the post-World War II sample period, suggesting that the business cycle may explain these deviations from the stationary random walk model. We directly measure the business cycle by the stochastically detrended level of industrial production. Regressing excess returns against this estimate of industrial production's cyclical component, our univariate test procedures provide little evidence that the resultant residuals exhibit persistence.

Of course, we cannot design a single univariate test which is most powerful against *all* alternatives. It is quite likely then that our χ^2_S and Bonferroni test statistics will not be able to detect potentially more complicated autocorrelation patterns present in the data once we have removed industrial production's cyclical component. That the default and term spreads can forecast these residuals implies that these market based variables provide information about excess returns beyond that contained in production's cyclical component and suggests that the business cycle itself does not explain all of the predictability in excess returns.

To more fully characterize the relationship between financial markets and the macroeconomy, this paper recognizes that macroeconomic time series are characterized by variable trends as well as recurrent cyclical fluctuations around this growth path. Our resultant multivariate

analysis concludes that while the business cycle contributes to the deviations from the stationary random walk model that we have documented for stock prices, predictable long term swings in expected returns arising from these variable trends remain.