Review of Financial Economics 24 (2015) 12–17

Contents lists available at ScienceDirect

Review of Financial Economics

journal homepage: www.elsevier.com/locate/rfe

Are equities good inflation hedges? A frequency domain perspective☆

Cetin Ciner *

Department of Economics and Finance, Cameron School of Business, University of North Carolina Wilmington, Wilmington, NC 28403, United States

ARTICLE INFO

Article history:
Received 28 April 2014
Received in revised form 5 December 2014
Accepted 10 December 2014
Available online 17 December 2014

JEL classification:
GO
G1

Keywords:
Inflation
Stock returns
Frequency domain

ABSTRACT

By using industry level data, we examine the relation between equity returns and inflation in a frequency dependent framework. Our analysis shows that a positive relation in fact exists between equity returns and high frequency inflation shocks for commodity and technology related industries. Since higher frequency shocks are independent from trend and are transitory in nature, our findings imply a positive relation between stock returns and the unexpected component of inflation. Furthermore, we show that the results are robust to firm-level data by using a sample from the oil industry. Hence, our study provides a new look at the impact of inflation on equities by showing the sensitivity of conclusions in prior work to frequency dependence in data.

1. Introduction

The relation between inflation and equity prices, motivated by the Fisher hypothesis, is one of the most frequently investigated topics in finance and economics. The underlying intuition of this hypothesis is that stocks represent claims on real assets and therefore, their valuations should increase with inflation. Thus, a positive relation is predicted between inflation and stock price movements, in which case it can be argued that equities provide a good hedge against the inflation risk.

Unfortunately, empirical studies, beginning with the early work of Bodie (1976) and Fama and Schwert (1977), consistently report a negative correlation between stock returns and inflation. Ang et al. (2011) and Hagmann and Lenz (2004) are examples of more recent papers that reach similar conclusions. One of the explanations offered for this counter-intuitive finding is the "proxy" hypothesis of Fama (1981) and Kaul (1987). The argument of these authors is that the documented negative correlation with inflation is spurious. It arises because the stock market anticipates the negative impact of higher inflation on growth, which lowers the market valuations.[1]

In this paper our objective is to provide a further examination of the relation between stock returns and inflation by using a recently developed frequency domain decomposition method by Ashley and Verbrugge (2009). The principal advantage of this method is to allow us to examine whether there is persistency dependence in the stock return-inflation linkage, which could occur if higher and lower frequency inflationary shocks have different effects on stock price movements.[2] We argue that differential effects of expected and unexpected inflationary shocks on stock valuations can be investigated by utilizing this method of analysis.

Specifically, lower frequency shocks are those that tend to be persistent and are likely to represent a continuation of the trend in inflation. Thus, at least to some extent, these shocks can be anticipated by market participants. Higher frequency shocks, on the other hand, are those with less persistence and are, therefore, transitory in nature. These shocks are difficult to forecast by definition and hence, represent the unexpected component of inflation. Hence, our empirical method provides a novel approach to define unexpected inflation as those corresponding to higher frequencies on the spectra. Our primary hypothesis is that if equities are good inflation hedges then we should expect a positive relation between stock returns and *unexpected* inflation.

Many studies in prior work have also examined the stock return–inflation linkage by obtaining proxies for expected and unexpected

☆ I am grateful to Richard Ashley for providing the code to conduct the frequency domain decomposition used in this paper. Comments provided by an anonymous referee significantly improved the paper. All remaining errors are mine.

* Tel.: +1 910 962 7497.
E-mail address: cinerc@uncw.edu.

[1] Geske and Roll (1983) and Pearce and Roley (1988) provide discussions of alternative hypotheses on the negative correlation between inflation and stock returns.

[2] While there are other frequency decomposition methods suggested in the literature, the Ashley–Verbrugge approach is unique because it is robust to feedback between the variables as discussed further below. In other words, it continues to be valid even when there is causality from equities to inflation, which is plausible under the present value models of asset prices.

http://dx.doi.org/10.1016/j.rfe.2014.12.001

inflationary shocks. For example, Fama and Schwert (1977) use changes in the short term interest rate to infer the expected inflation. Bodie (1976) relies on ARIMA models to construct proxies for expected inflation. Hagmann and Lenz (2004) and Hess and Lee (1999) use vector autoregression (VAR) models to decompose inflation into expected and unexpected components. In related work, McQueen and Roley (1993) and Pearce and Roley (1988) rely on forecasts by Money Market Services International to identify the surprise element in inflation announcements. More recently, Wei (2009) uses a time series regression model that also includes lagged values of monthly unemployment rate to estimate the unexpected component of inflation. These studies in general support the conclusion of earlier papers that the relation between stock returns and unexpected inflation is also negative.

Furthermore, in papers closely related to the present article, Lee and Ni (1996) use the Chebyshev filter to obtain estimates of unexpected inflation as the transitory component of inflationary shocks. These authors argue that there are differences in the relation between stock returns and inflation based on the frequency components of inflation. Kim and In (2006) rely on a multi-scale wavelet decomposition approach to examine the same relation. Interestingly, they detect a positive relation between inflation and stock returns at the shortest scale, which corresponds to the highest frequency shocks. We provide a further discussion in the following section on why we have preferred the Ashley–Verbrugge decomposition in the present paper rather than the abovementioned method of analyses.[3]

We also examine the stock return–inflation relation by relying on industry portfolios. This permits us to determine whether some sectors of the stock market may be considered better hedges against inflation. For example, Boudoukh et al. (1994) argue that non-cyclical stocks tend to be more positively correlated with inflation than cyclical stocks. Stocks of natural resource companies are usually considered as better inflation hedges because of the sensitivity of commodity prices to inflation. In the empirical analysis, therefore, we first calculate the conventional inflation betas using 48 industry portfolios by estimating conventional stock return–inflation regressions. Our data cover the period between 1990 and 2013 and hence, a further contribution of our analysis is to update evidence in the literature for a more recent period. Consistent with prior work, we find that, in general, inflation betas are negative.

Subsequently, we estimate the same model by replacing inflation with its frequency components obtained by the Ashley and Verbrugge (2009) decomposition. The results of this analysis lend support to the central claim of the paper, which is that the stock return–inflation relation shows dependency on the persistence of inflationary shocks. Specifically, we find that while long term, trend shocks, replicate the negative inflation betas obtained above as expected, inflation betas for unexpected (high frequency) shocks are in fact positive for 18 industry portfolios. These positive unexpected inflation betas exist in commodity-sensitive (such as coal, mines, oil, gold and agricultural) and technology-related industries (such as telecoms, software and chips) among others.

The first part of the empirical analysis utilizes value-weighted industry portfolios and hence, could simply be an artifact of data, since larger companies dominate value-weighted indices. To investigate the robustness of the findings to firm size, we conduct the analysis by using equally-weighted industry portfolios. We find largely similar results. Again, unexpected inflation betas for commodity-sensitive industries (gold, mines and oil) as well as technology-related industries (telecom, hardware and chips) have positive, and statistically significant, unexpected inflation betas.

In the next section, we outline the statistical method of analysis used in the paper. In Section 3, we present the data set, and discuss the empirical findings in Section 4. We provide the concluding comments of the paper in the final section of the study.

[3] We thank an anonymous referee for recommending this discussion.

2. Statistical method of analysis

Conventional time domain regressions are linear and force a fixed coefficient to describe the relation between the variables that is supposed to be constant at all frequencies. For instance, in the context of our paper, the conventional regression analysis suggests that the sensitivity of industry portfolio returns to transitory (high frequency) shocks in inflation is exactly the same as it is to permanent (low frequency) shocks. However, there is no a priori reason to expect this to always obtain and researchers have long recognized that estimating a frequency dependent regression model, in which the coefficient is permitted to vary over time, is likely to yield richer dynamics. Early research by Hannan (1963) and Engle (1974), further developed by Tan and Ashley (1999), suggests to transform the time series regressions into frequency domain by means of spectral regression models. For example, consider the following generalized linear regression model:

$$Y = X\beta + \epsilon, \epsilon \sim N\left(0, \sigma^2 \mathrm{I}\right) \tag{1}$$

in which Y is T × 1 and X is T × K. The objective of the approach is to transform the equation in such a manner that the components of the variables correspond to frequencies rather than time periods. This is accomplished by pre-multiplying the regression by a T × T matrix A, whose (s,t)th element is given by

$$A_{s,t} = \begin{cases} \left(\frac{1}{T}\right)^{1/2}, & \text{for } s = 1; \\ \left(\frac{2}{T}\right)^{1/2} \cos\left[\frac{\pi s(t-1)}{T}\right], & \text{for } s = 2, 4, 6, \ldots, (T-2) \text{ or } (T-1) \\ \left(\frac{2}{T}\right)^{\frac{1}{2}} \sin\left[\frac{\pi(s-1)(t-1)}{T}\right], & \text{for } s = 3, 5, 7, \ldots, (T-1) \text{ or } T; \\ \left(\frac{1}{T}\right)^{\frac{1}{2}}(-1)^{t+1}, & \text{for } s = T \text{ when } T \text{ is even} \end{cases}$$

$$AY = AX\beta + A\epsilon$$

A is an orthogonal matrix and the pre-multiplication gives:

$$Y^* = X^*\beta + \epsilon^*, \epsilon^* \sim n\left(0, \sigma^2 \mathrm{I}\right). \tag{2}$$

In Eq. (2), the components of the variables now represent frequencies instead of time periods. Next, the T frequency components are partitioned into M frequency bands and dummy variables are created to define M vectors of length T, which can be written as D^{*1}, ….. D^{*M}. These dummy variables are used in a manner that for elements, which fall into the sth frequency band, D^{*s} equals X_j^* and the elements are zero otherwise. Consequently, the regression equation can be rewritten as follows:

$$Y^* = X^*_{\{j\}}\beta_{\{j\}} + \Sigma_{m=1}^{M}\beta_{j,m}D^{*m} + \epsilon^* \tag{3}$$

in which $X^*_{\{j\}}$ is the X^* with its jth column deleted and $\beta_{\{j\}}$ is the β vector with its jth component deleted. Hence, frequency-dependent coefficients $\beta_{j,1}$ … … .. $\beta_{j,M}$ can be estimated and hypotheses tests can be conducted on the significance of these parameters.

However, Ashley and Verbrugge (2009) argue that an important weakness exists in the above approach. In particular, since pre-multiplying with matrix A mixes past and futures values of the variables, the M frequency components will be correlated with the error terms, if there is feedback between the dependent and any of the independent variables, which is likely in finance and economics data. This would yield inconsistent estimates if the partitioned frequency components are used in an OLS regression.

The main contribution of Ashley and Verbrugge (2009) is that they present a solution to this problem by applying a one-sided, rather