

Olesen, Jan Overgaard

Working Paper

Stocks hedge against inflation in the long run: evidence from a cointegration analysis for Denmark

Working paper, No. 6-2000

Provided in Cooperation with:

Department of Economics, Copenhagen Business School (CBS)

Suggested Citation: Olesen, Jan Overgaard (2000) : Stocks hedge against inflation in the long run: evidence from a cointegration analysis for Denmark, Working paper, No. 6-2000, Copenhagen Business School (CBS), Department of Economics, Frederiksberg, <https://hdl.handle.net/10398/7557>

This Version is available at:

<https://hdl.handle.net/10419/208428>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

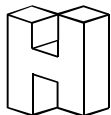
Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc-nd/3.0/>



Institut for Nationaløkonomi

Handelshøjskolen i København

Working paper 6-2000

STOCKS HEDGE AGAINST INFLATION IN THE LONG RUN: EVIDENCE FROM A COIN- TEGRATION ANALYSIS FOR DENMARK

Jan Overgaard Olesen

This version, March 2000
First draft, August 1998

Stocks Hedge against Inflation in the Long Run: Evidence from a Cointegration Analysis for Denmark^{*}

by

Jan Overgaard Olesen

Department of Economics and EPRU[§]
Copenhagen Business School
Denmark

Abstract

We suggest an alternative approach to testing whether stocks provide a hedge against inflation in the long run. Based on a simple structural model, we test the hedge hypothesis in terms of the long-run linkage between stock prices and the general price level, as estimated by cointegration analysis. Using data for the Danish stock market over the post-World War II-period, results give strong support for the hedge property, defined in the narrow sense of a perfect hedge. This contrasts with the weak support found in the literature and also represents stronger support than produced by standard methods. We argue that our approach has the advantage of allowing for a clear distinction between short- and long-run dynamics of stock prices which adjust slowly to long-run equilibrium.

Postal address: Department of Economics, Copenhagen Business School, Solbjerg Plads 3 (5th), DK-2000 Frederiksberg, DENMARK. Phone: +45 3815 2575. As of March 20 2000: Danmarks Nationalbank, Havnegade 5, DK-1093 Copenhagen K, DENMARK. Phone: +45 3363 6363.

^{*} We have benefited from comments by participants at the EPRU workshop “The Stock Market and The Macroeconomy”, Copenhagen Business School, May 1998. In particular, we thank Ole Risager for comments.

[§] The Economic Policy Research Unit (EPRU) is financed by a grant from The Danish National Research Foundation.

1. Introduction

Stocks are said to provide a hedge against inflation if they compensate investors completely (and not by more) for increases in the general price level through corresponding increases in nominal stock returns, thereby leaving real returns unaffected. That is, stocks hedge against inflation if their real value or purchasing power is immune to changes in the general price level.

Whether or not stocks hedge against inflation is relevant to any rational investor who cares about real wealth. The above definition is one of a perfect hedge as it demands a one-for-one compensation for inflation. This contrasts with the weaker notion of an imperfect (or partial) hedge, as often encountered in the literature, which requires the relation between nominal stock returns and inflation (or equivalently, between nominal stock prices and the general price level) to be significant and positive but it may be less or larger than one-for-one. However, with an imperfect hedge, the real value of a portfolio of stocks is subject to uncertainty due to the uncertainty about future inflation. This is not the case when the hedge is perfect. As we interpret an inflation hedge as a device of eliminating the uncertainty deriving from inflation uncertainty, we shall throughout use the term in its most restrictive sense of a perfect hedge.

Apriori it can be argued that stocks should provide a hedge against inflation, at least in the long run where firms' profit margins can reasonably be assumed to be fixed. The argument is that stocks are claims on current and future profit opportunities which in the long run (with profit margins being fixed) increase with the general price level in relation one-for-one, that is, in the long run stocks are basically claims on real profit opportunities. As a result, we should expect the real value of stocks to remain unaffected by inflation and, hence, stocks should hedge against inflation in the long run. What happens in the short run is, on the other hand, more ambiguous because slow adjustment in output prices and real production imply that profit margins may be significantly affected by inflation.

Whether stocks also provide a hedge against inflation empirically has been studied extensively in the literature, see e.g. Fama and Schwert (1977), Gultekin (1983), Boudoukh and

Richardson (1993), Ely and Robinson (1997) and Barnes *et al.* (1999). With the only exception of Ely and Robinson (1997), cf. below, the literature has based its inference on return regressions where nominal stock returns are regressed on inflation and possibly further explanatory variables such as real production growth and changes in a relevant discount rate measure. The inflation hedge hypothesis is then put to a test by testing whether the coefficient to inflation is significant and equal to 1¹. Results of the literature are fairly mixed, but a general conclusion is that stocks do not hedge against inflation in the short run (investment horizons less than 1-2 years), where inflation usually turns out to have an insignificant effect on stock returns. In fact, at short horizons the estimated relation between nominal stock returns and inflation may even be negative, see e.g. Fama and Schwert (1977) and Gultekin (1983). There is some evidence of a significant positive relationship on longer horizons (more than 2 years) but often with a coefficient different from 1 so that the inflation hedge is not perfect, cf. Boudoukh and Richardson (1993). Hence, the hedge hypothesis comes closer to receiving support at longer horizons but the evidence is still weak. On balance it therefore seems that the empirical evidence tends to reject the hypothesis of stocks providing a (perfect) hedge against inflation.

This paper tests the inflation hedge hypothesis for stocks by taking a different approach to that used in the literature. We test the hypothesis by focusing on the long-run relation between stock prices and the general price level rather than the relation between stock returns and inflation. Most importantly, this shift of focus allows us to take account of slow adjustment in stock prices in the event of inflation. The latter is from the outset precluded in the standard return regressions approach which (implicitly) assumes that stock prices adjust completely to inflationary shocks over the prespecified, fixed investment horizon, see section 5 below for a further discussion. We focus explicitly on the long-run horizon where the fixed-profit-margin assumption underlying the hedge hypothesis *apriori* seems most relevant. We

¹ Some studies frame the test in terms of real rather than nominal stock returns, testing whether inflation has a significant influence on real stock returns, see for instance Fama (1981) and Kaul (1987). A survey of the literature including a detailed account of the empirical results is provided by Frennberg and Hansson (1993). The latter study at the same time represents an exception in the literature as the authors conclude that Swedish stocks provide a hedge against inflation even at fairly short horizons (down to one month). Another survey of the literature can be found in Sellin (1998). He concludes that “Stocks seem to be a good hedge against both expected and unexpected inflation at longer horizons” (Sellin 1998, p. 25). However, this

proceed as follows. Motivated by a simple theoretical framework, we formulate a structural model for stock prices which includes the general price level, real production and stock investors' discount rate as explanatory variables. We identify the long-run relationships between the variables by cointegration analysis, using the cointegrated VAR-model, see e.g. Johansen (1996). We estimate a cointegrating relation for stock prices and, finally, test the inflation hedge hypothesis by testing whether this relation implies a one-for-one relationship between stock prices and the general price level.

We test the hypothesis for the market portfolio of Danish stocks, using annual data from 1948 to 1996. While the sample may be considered small in terms of the number of observations, the sample period spans many years which is crucial for the analysis of "the long run". In the empirical analysis, we use small sample versions of tests whenever possible. Moreover, we check the robustness of results from the cointegrated VAR model by also using single-equation-cointegration-methods to test the hedge hypothesis.

Our approach has similarities with that of Ely and Robinson (1997) who also differ from the standard literature by focusing on the relation between stock prices and the general price level in testing the inflation hedge hypothesis. Ely and Robinson (1997) test the hypothesis for 16 OECD countries, based on impulse response analysis in a cointegrated VAR model with 4 variables - stock prices, the general price level, real production and money supply. They find for almost all countries that stocks overcompensate for inflation and conclude, using an imperfect hedge definition, that stocks hedge against inflation. However, using the more restrictive definition of a perfect hedge, the evidence in Ely and Robinson (1997) does not give support to the hedge hypothesis.

Our approach differs from Ely and Robinson (1997) in several ways. First of all, we differ in the definition of an inflation hedge. In addition to the use of a perfect rather than an imperfect hedge definition, we define an inflation hedge in terms of the 'partial' sensitivity of stock prices wrt. the general price level within the context of a structural model for the former. Thus, we address the question: What happens to stock prices in the event of shocks to the

conclusion is based on an imperfect hedge definition, which allows stock prices (or returns) to respond more

price level, all other factors (real production and the discount rate) kept constant ? Ely and Robinson (1997), on the other hand, examine the response in stock prices within a VAR model which we interpret as a reduced form model for stock prices and the price level where real production and the money stock are the ‘driving’ (exogenous) variables². Hence, they address the question: What happens to stock prices in the event of shocks to the price level, when other factors (e.g. real production and the discount rate) are allowed to vary ? Our *ceteris paribus* definition of an inflation hedge resembles that used in the literature of return regressions.

Second, we test the hedge hypothesis in terms of a cointegrating relation for stock prices and, hence, do not rely on impulse response analysis as in Ely and Robinson (1997). This may be viewed as an advantage, given the critique raised by e.g. Faust and Leeper (1997), who show that results from impulse response analysis depend crucially on the assumptions needed to identify the underlying structural shocks of the VAR model. This may question the robustness of results derived from impulse response analysis. Moreover, by focusing on the cointegrating relation, we can perform an explicit parametric test of the hedge hypothesis instead of the ‘qualitative’ test criteria used in Ely and Robinson (1997)³.

Finally, we can test whether the underlying framework of our approach - the structural model for stock prices - is reasonable empirically by testing whether it is validated as a cointegrating relation. This turns out to be the case, implying that we can have (some) confidence in the framework underlying the test of the hedge hypothesis. For instance, the evidence of cointegration suggests that we do not lack an important variable in modeling the long-run linkages between stock prices and the general price level. Such a validity test of the underlying framework is not (directly) possible in the approach of Ely and Robinson (1997).

than proportionately to shocks to the general price level (or to inflation).

² Ely and Robinson (1997) do not provide a theoretical foundation for their VAR model.

³ Based on the impulse response analysis, Ely and Robinson (1997) test the hedge hypothesis at a qualitative level, concluding that “In those cases where the impact on stock prices is significantly positive (negative) and/or where the impact on goods prices is significantly negative (positive), stocks offer (do not offer) a hedge against inflation in the sense that the relative value of stock prices to goods prices rises (falls)” and “Stocks can also be said to offer a hedge in those cases where neither stock price nor goods price innovations are statistically significant”, Ely and Robinson (1997, page 151).

Compared to the existing literature, the contribution of the paper is three-fold. First of all, we suggest an alternative approach to testing the inflation hedge hypothesis. Second, it turns out that results give strong support to the hypothesis which contrasts with the weak support found in the literature. Third, the paper provides results for Denmark, a case which to our knowledge has not been examined thoroughly before⁴.

The paper is organized as follows. In section 2 an operational empirical model for the long-run is formulated. Section 3 reviews the data and section 4 reports the empirical results. Section 5 concludes the paper with a summary and a comparison of our approach with that used in the literature.

2. An Empirical Model for the Long Run

We formulate an empirical structural model for stock prices based on a simple theoretical framework that links stock prices to the general price level. The framework is *ad hoc* and rests on a set of assumptions which are restrictive but facilitate the formulation of an empirically tractable model. We focus on the long-run horizon with the objective of a model that can act as a good approximation to the long-run movements in stock prices. This provides us with a sound empirical (and a theoretical) foundation for testing the inflation hedge hypothesis in the long run. Whether the model actually is a good approximation, is tested as part of the empirical analysis by testing whether it can be validated as a cointegrating relation for stock prices.

The starting point is the usual 1-period no-arbitrage relation between stocks and bonds under the assumption of perfect capital markets. Excluding risk premia, this relation demands that the expected 1-period holding return on stocks, consisting of a capital gain and a dividend yield, is equal to the 1-period return (yield-to-maturity) on bonds:

⁴ Bonnichsen (1983) is an informal study of the relationship between Danish stock returns and inflation in the period 1900-1982. He examines whether the nominal stock return exceeds inflation at long investment horizons, that is, whether the real return at long horizons is positive, and concludes this to be the case. However, this evidence does not address the basic issue whether stocks *hedge* against inflation. The latter requires an analysis of how stock returns (or stock prices) *respond* to *changes* in the inflation rate (or the general price level). Thus, *a priori* the real return on stocks may still be positive in a situation where the nominal stock return *does not* respond to changes in the inflation rate, that is, in a situation where stocks *do not* hedge against inflation.

$$(1) \quad \frac{Q_{t+1}^e - Q_t}{Q_t} + \frac{D_{t+1}^e}{Q_t} = B_t$$

where Q_t is the (ex dividend) stock price per share at time t , D_{t+1} is the dividend payment per share during period $t+1$ and B_t is the 1-period bond return as of time t . Superscript “e” denotes expectations on unknown future variables. The stock is assumed to be a claim on a representative firm (in our case representative for all firms listed at the Copenhagen Stock Exchange).

We shall assume that investors only form point expectations on future variables, i.e., that ‘Certainty Equivalence’ applies, and that investors, furthermore, expect bond returns to be constant over time. This, and the exclusion of rational bubbles, gives the forward-looking stock price solution⁵:

$$(2) \quad Q_t = \sum_{i=0}^{\infty} \left(\frac{1}{1+B_t} \right)^{i+1} D_{t+1+i}^e$$

which determines the stock price as the expected discounted value of all future dividend payments.

Now make the following assumptions:

(A1) Constant profit margin π^* , i.e., profits $\Pi_t = \mathbf{p}^* P_t Y_t$

(A2) Output price P_t and real production Y_t are expected to grow at constant growth rates g_p and g_y , respectively, i.e.

$$P_T^e = P_t (1 + g_p)^{(T-t)} \quad \text{and} \quad Y_T^e = Y_t (1 + g_y)^{(T-t)} \quad (T \geq t)$$

(A3) All profits are paid out as dividends each period, i.e. $D_t = \Pi_t$

P_t is the price of the firm's product, Y_t is the production of the same product and Π_t denotes total 'profits' of the firm, assumed to be a constant fraction π^* (the profit margin) of the value of production. Π_t should be interpreted as the earnings that the firm generates to stock holders and could, to be specific, be defined as the value of production (value-added) less labor costs, accounting for 'pure' profits in a firm without capital, respectively, 'pure' profits plus capital rent in a firm with capital. By assuming a constant profit margin, our basic working hypothesis is that any fluctuations in the profit margin are purely short-run (business-cycle) phenomena which are eliminated in the long run. In particular, we assume that any changes in the relative prices between output and inputs (e.g. real wages and real oil prices) are either reversed or validated by average productivity changes in the long run. In theory, the assumption of a constant profit margin will, for instance, hold for a perfectly competitive firm with a Cobb-Douglas production technology, in which case the capital income share (taken to be the profit margin) is fixed. The assumption may also be justified (for the long term) by empirical observations, as evidence suggests that the aggregate profit share in the Danish economy has been fairly stable over the period since World War II⁶. g_p and g_y , finally, denote the expected inflation rate and the expected real growth rate. We assume that the number of shares in the firm is constant over time and normalize it to 1.

From (A3) and the expected dynamics of profits implied by (A1) and (A2), the solution for stock prices becomes:

$$(3) \quad Q_t = \frac{\Pi_t}{R_t} = \frac{p^* P_t Y_t}{R_t}$$

where

⁵ Assuming that the forward-looking stock price solution exists, i.e. that dividend payments are expected to grow at a rate less than B_t .

⁶ A possible measure of the aggregate profit share is the ratio of Gross Operating Surplus to GDP at factor cost (both in current prices), as defined by the National Account Statistics. Using annual data for the private sector over the period from 1948 to 1996, this ratio has varied within the narrow interval between its low of 38% in 1980 and its high of 49% in 1951.

$$(4) \quad R_t \equiv \frac{1 + B_t}{(1 + g_p)(1 + g_y)} - 1 \approx B_t - g_p - g_y \text{ (for } g_p \text{ and } g_y \text{ "small")}$$

R_t is the (*ex ante*) growth-adjusted real discount rate, defined as the nominal bond return adjusted for expected inflation and expected real growth. (3) is basically just a variant of the standard Gordon-growth-formula for the price of a stock with a constant discount rate and constant dividend growth, cf. e.g. Campbell *et al.* (1997). (3) only differs by allowing for time-variation in the discount rate and by having replaced dividends by profits.

We shall say that stocks provide a hedge against inflation if shocks to the general price level result in proportional changes in stock prices when controlling for other relevant factors. Our simple framework highlights why we should expect stocks to hedge against inflation in the long run. Thus, consider a shock to current prices P_t , reflecting the outcome of past inflation. Such a shock translates *ceteris paribus* into a proportional change in the value of production ($P_t Y_t$) and - due to the constant profit margin - profits (\tilde{O}_t). Because prices and production are expected to grow over time at fixed (unaffected) rates, expected future profits, likewise, change proportionally. As a result, current stock prices (Q_t) change proportionally, confirming the hedge property, cf. also (3). Note the *ceteris paribus* (or partial) content of the hedge property as real production and the discount rate are held fixed in the argument⁷. Based on (3), we formulate the following empirical model expressed in logarithmic terms (lower case letters denote corresponding log-levels):

$$(5) \quad q_t = b_0 + b_1 p_t + b_2 y_t + b_3 r_t + e_t$$

The ϑ_i 's are coefficients (including a constant term) to be estimated and M_t is the residual of the equation.

⁷ The literature using the return regressions approach also controls for other relevant factors by regressing stock returns not only on inflation but also on further explanatory variables (e.g. the real growth rate), and focusing on the direct effect from inflation in testing the hedge hypothesis. From an econometric point of view, the inclusion of other relevant factors is important in order to avoid an omitted-variables bias in the estimate of the inflation effect. The latter is also true in our approach. We test (indirectly) for having omitted

(5) explains the long-run movements in stock prices by the long-run movements in the general price level, real production and the real discount rate. As the variables considered are non-stationary, cf. section 3, we have to use cointegration techniques in estimating (5). If our framework is valid empirically, we should expect (5) to be a cointegrating relation, i.e., a stable, long-run equilibrium relation for stock prices⁸. Whether this is actually the case, is tested as an initial step of the empirical analysis. On a validation of (5), we can then test the inflation hedge hypothesis. Given our definition of the hedge property, a formal test of the hypothesis can be framed in terms of the coefficient to the general price level, ϑ_1 , measuring the direct (or partial) effect of the price level on stock prices. The hedge hypothesis stipulates that there exists a long-run linkage between stock prices and the price level, i.e., that the price level is significant in (5) ($\vartheta_1 \neq 0$) and, moreover, that the elasticity of stock prices wrt. the price level is exactly one ($\vartheta_1=1$). Hence, the hypothesis is supported if, and only if, the estimated ϑ_1 is significant and, furthermore, not significantly different from one.

3. The Data

All data are annual and cover the period 1948-1996. The source database is Nielsen, Olesen and Risager (1997).

Stock prices are measured by the overall stock price index by Statistics Denmark, comprising all firms listed at the Copenhagen Stock Exchange (CSE). For the general price level, we use the official Consumer Price Index (CPI). We consider the question of whether stocks hedge against inflation to be most interesting in terms of CPI inflation because stock investors - ultimately being consumers - care about real wealth in terms of consumption bundles.

Moreover, CPI is the price measure encountered in the literature. We choose to proxy real production by a deterministic trend. This may be justified by the fact that we are interested in modeling the movements in production over long horizons and, for this purpose, a trend may be a reasonably good approximation⁹.

variables important for the long-run modeling of stock prices, by testing whether the model provides a cointegrating relation.

⁸ Formally, we have cointegration if, and only if, the residual term M_t is stationary. In this case, (5) serves as an 'attractor' for the included variables, see e.g. Engle and Granger (1991) for an interpretation of the concept of cointegration.

⁹ What we need is a proxy for real production which results in (5) being a cointegrating relation. This turns out to be the case when using a deterministic trend. We have tried several explicit production measures (e.g. real GDP for the overall economy, for the private sector and for manufacturing) but without any further

In order to estimate (5), we also need a proxy for the unobservable (*ex ante*) growth-adjusted real discount rate r_t . We use a discount rate measure which in a given year is calculated as the difference between the yield-to-maturity of a 10-year government bond and the historical inflation rate over the 5-year period preceding that year. This proxy results - compared to other discount rate proxies that we have examined - in the strongest evidence that (5) is a cointegrating relation. We consider this to be a valid criterion for choosing the proxy because we only want to formulate an empirically valid framework prior to testing the inflation hedge hypothesis, i.e., to formulate a model that captures the important long-run features of stock prices. The latter is evidenced by the presence of cointegration. The stronger cointegration could in fact be interpreted as evidence that this proxy is particularly good at modeling the long-run movements in the ‘true’ discount rate.

Notice that the use of proxies introduces measurement errors in the explanatory variables in (5). However, as long as the measurement errors are stationary, this does not affect the inference on the cointegrating relation, cf. Hamilton (1994).

Figure 1 shows the data.

< Insert Figure 1 around here >

< Insert Tables 1.a and 1.b around here >

Unit root tests are performed using both the Phillips and Perron (1988) Z_t -test (PP) and the test by Kwiatkowski *et al.* (1992) (KPSS), cf. Tables 1.a and 1.b¹⁰. Using conventional significance levels, both tests clearly support that stock prices and the discount rate proxy are integrated of order 1 (I(1)), i.e., are non-stationary with stationary first differences. For CPI, the PP test concludes integration of (at least) order 2 (I(2)), i.e., both levels and first

success in establishing a cointegrating relation, cf. Appendix A, which reports the results of estimating alternative candidates for a cointegrating relation for stock prices, using alternative measures of both production and the general price level.

¹⁰ We have used a maximum of 6 lags in both tests because the test statistics become reasonably stable within this lag length. The evidence in Kwiatkowski *et al.* (1992) also suggests that, for our sample size, the KPSS test has a reasonable size and power at a lag length around 4 to 6.

differences are non-stationary, using a strict 5% significance level. However, the PP test consistently supports the I(1) hypothesis at a 10% level. The KPSS test strongly supports the I(1) hypothesis for CPI when allowing for serial correlation in the disturbance term (lag length $l^3 1$). Overall, the evidence is therefore in favor of I(1).

To conclude, all series are I(1). The exclusion of the possibility of I(2)-behavior means that the standard Johansen-procedure can be used for estimating (5). Moreover, as (5) is balanced in terms of unit root behavior, single-equation-cointegration techniques can be used for estimation purposes.

4. The Empirical Results

Motivated by (5), we formulate a VAR model using stock prices, the general price level (CPI) and the discount rate proxy as the endogenous variables, and including a deterministic trend. To outline the model, let the endogenous variables be described by the column vector $X_t = (q_t, p_t, r_t)'$. Following the notation of Johansen (1996), the VAR model can be written in its reduced vector error-correction form (VECM) as:

$$(6) \quad \Delta X_t = \Pi X_{t-1} + \sum_{i=1}^{k-1} \Gamma_i \Delta X_{t-i} + \Phi D_t + e_t$$

where k denotes the lag length, Π and Γ_i are matrices of dimensions 3×3 and D_t is a 2×1 vector containing the deterministic terms. We allow for a constant term and a deterministic trend, i.e., $D_t = (1, t)'$. F is the 3×2 matrix which contains the coefficients to the deterministic terms. e_t is the vector of disturbance terms, assumed to be identically distributed “white noise”.

The rank of matrix Π , denoted by r , determines the number of cointegrating relations among the three endogenous variables. If Π has zero rank ($r=0$), there is no cointegration in the data and (6) becomes a VAR model in first differences only because the level term disappears. If Π has a non-zero, but reduced rank ($0 < r < 3$), (6) is a cointegrated VAR model with r (linearly

independent) cointegrating relations. In this case, Π can be written as the product of two full column rank matrices α and β of dimensions $3 \times r$, i.e., $\Pi = \alpha\beta'$, and (6) can be rewritten as:

$$(7) \quad \Delta X_t = \alpha\beta' X_{t-1} + \sum_{i=1}^{k-1} \Gamma_i \Delta X_{t-i} + \Phi D_t + e_t$$

Each column vector in the β -matrix corresponds to a cointegrating relation in the sense that the linear combination $\beta_i' X_t$, where β_i (here) denotes the i 'th column vector of β , is stationary. $\beta_i' X_t$ corresponds to the usual error-correction-term in single-equation cointegration analysis. Each vector is called a cointegrating vector and there exists a total of r (linearly independent) cointegrating vectors. The matrix α contains the adjustment coefficients by which each cointegrating relation affects the short-run dynamics of the endogenous variables. For example, element α_{ji} in α captures by how much the short-run dynamics of variable j in X_t (ΔX_{jt}) responds to the equilibrium error in cointegrating relation no. i ($\beta_i' X_t$). Finally, if Π has full rank ($r=3$), we have in principle 3 cointegrating relations, which is only possible if all the variables are stationary.

We restrict the deterministic trend to be in the cointegrating space, precluding the possibility of a quadratic trend in the endogenous variables, cf. Johansen (1996). The latter assumption seems both plausible and, at an informal level, cf. Figure 1, validated by the data. The estimation is therefore based on the VAR specification:

$$(8) \quad \Delta X_t = \alpha\beta^{*'} X_{t-1}^{*} + \sum_{i=1}^{k-1} \Gamma_i \Delta X_{t-i} + m_0 + e_t$$

where $\beta^{*'} = (\beta', \rho_1)'$ and $X_t^{*'} = (X_t', t)'$, i.e., the trend is included as part of the cointegration term. m_0 is the vector of unrestricted constant terms while the $r \times 1$ vector ρ_1 contains the coefficients to the trend in the cointegrating relations. In the empirical analysis, interest focuses on, first of all whether there exists any cointegrating relations or vectors β^* , and, secondly, on the coefficients of the cointegrating vectors, β^* , in particular, the coefficient to the general price level.

As the initial step in the estimation, the appropriate lag length (k) of the VAR model has to be determined¹¹. Various procedures can be used, including the explicit testing on lag coefficients in a “general-to-specific” procedure and the use of information criteria. Using the “general-to-specific” procedure, we start out with 6 lags which is sufficient to ensure that the white noise requirements on the disturbance term are fulfilled. We then successively remove insignificant lags from the top, performing a Likelihood Ratio test of the hypothesis that all coefficients at the largest lag are zero¹². This procedure results in a lag length of $k=4$, using conventional significance levels. The test for removing all variables at lag 4 leads to a clear rejection (critical significance level of 0.2%), while the hypothesis of reducing the lag length from 5 to 4 is firmly accepted (critical significance level of 58%). A lag length of 4 is supported by the Hannan-Quinn and Akaike information criteria while the Schwarz criterion suggests a shorter lag length of 2.

Table 2 reports both univariate and multivariate specification tests of the VAR model with 4 lags. Diagnostics for each equation in the model, including fitted values for the endogenous variables, are furthermore graphed in Figure 2. The specification tests test whether the residuals from the VAR model fulfill the white noise requirements of being serially uncorrelated, homoskedastic and normally distributed. According to the univariate test, the hypothesis of normally distributed residuals is rejected for the discount rate equation, using conventional significance levels. For the price level equation, the normality hypothesis is close to a rejection. However, the normality assumption is not crucial to the cointegrated VAR model, see Johansen (1996, Part II), who shows that it is a sufficient condition for using this method that the disturbance terms are identically distributed over time. The violation of the normality hypothesis is therefore not a problem for the inference to be drawn. There are no signs of misspecification according to the other, more critical specification tests for serial correlation and heteroskedasticity. Hence, we conclude that the VAR model with 4 lags is well specified and proceed with this specification.

¹¹ Estimations are performed in PCFIML, cf. Doornik and Hendry (1997).

¹² We use the approximate F-form of the Likelihood Ratio test suggested by Rao, cf. Doornik and Hendry (1997). This F-form which corrects for degrees of freedom is generally considered to have better small sample properties than the uncorrected χ^2 -form.

< Table 2 >

< Figure 2 >

The cointegration part of the VAR model (I and ϕ^*) is estimated by Maximum Likelihood, using the Johansen procedure, cf. e.g. Johansen (1996). Table 3 shows the (standardized) estimates of α and β^* together with estimated eigenvalues. Table 3 also reports statistics from trace tests on the rank of Π . Two trace test statistics are shown. The first statistic which is the one used in Johansen (1996) is the outcome of an asymptotic test. The evidence in Reimers (1992) suggests that this test is “over-sized” in small samples, implying that when using this test we tend to accept too many cointegrating relations compared to the significance level which we are actually willing to use. Based on this evidence and the fact that we have to deal with a small sample, we have more faith in the second trace test which adjusts the former test for degrees of freedom in the way discussed by Reimers (1992). This test is reported to have significantly better small sample properties in the sense that the actual significance levels of the test come close (closer) to the nominal levels in small samples.

< Table 3 >

Both rank tests lead to the conclusion that there is at least one cointegrating relation as both tests firmly reject the hypothesis of no cointegration ($r=0$) at conventional significance levels. The first (asymptotic) trace test also rejects the hypothesis of 1 cointegrating relation in favor of the alternative of more than 1 cointegrating relation. However, this hypothesis cannot be rejected according to the second (degrees-of-freedom-adjusted) trace test. Based on the latter test, we conclude that there is one and only one cointegrating relation between the variables ($r=1$). The second trace test gives a clear rejection of the hypothesis of no cointegration (the critical significance level is 1.9% by linear interpolation). Hence, the evidence of cointegration is strong.

The econometric identification of the cointegrating relations is relatively straightforward with only 1 cointegrating relation because normalizing on one of the variables suffices. Motivated by the modeling framework of section 2 (and the lack of an obvious alternative), we interpret

the cointegrating relation as a model for stock prices and normalize on this variable. The resulting estimates of the normalized cointegrating vector and the corresponding adjustment coefficients appear in Table 3 as the first column of β^* (i.e., β_1^*), respectively, the first column of α (i.e., the adjustment coefficients wrt. $\vartheta_1^* X_t^*$). The assumption that the cointegrating relation is a model for stock prices is actually supported by the estimates of the I-coefficients, because the error-correction in the short-run dynamics is strong in the direction of stock prices, whereas the corrections in the directions of the price level and the discount rate are very small in magnitude and can actually be shown to be insignificant, cf. below. The estimation gives the following long-run model for stock prices (indicative standard errors of the parameter estimates in parenthesis)¹³:

$$(9) \quad q_t = 0.96 + \underset{(0.13)}{1.04} p_t + \underset{(0.009)}{0.011} t - \underset{(2.4)}{5.42} r_t$$

All coefficients have signs consistent with theory. The trend may appear to be insignificant, using the indicative standard error, but we proceed with (9) because our interest lies with the price level coefficient and we do not want to condition the inference on the coefficients to the remaining variables.

We take the estimated cointegrating relation as evidence in favor of the modeling framework of section 2, hence establishing a firm empirical framework within which to test the inflation hedge hypothesis. The hedge hypothesis is tested by Likelihood Ratio (LR) tests on the coefficient to the price level in the cointegrating relation. These tests compare the likelihood of the unrestricted VAR model (where the price level coefficient can vary freely) with the likelihood of the restricted VAR model (where the price level coefficient is restricted). Testing, first, the null hypothesis that the price level has an insignificant effect on stock prices ($\vartheta_1=0$), the outcome is a LR test statistic of 11.8 which has to be compared with a $\chi^2(1)$ -distribution. The critical significance level is for all practical purposes zero, leading to a strong rejection of the null. Hence, the price level has a significant effect on stock prices in

¹³ The constant term in (9) is calculated from the formula $r_0 = (a' a)^{-1} a' m_0$ where μ_0 is the unrestricted constant term, cf. (8), and p_0 is the component of this constant term which enters the cointegrating relation, see Johansen (1996, p. 81). I here denotes the first column of the estimated I-matrix in Table 3.

the long run. Next, testing the null hypothesis that stock prices and the price level move one-for-one ($\vartheta_1=1$) gives a test statistic of 0.04 which, again, has to be compared to a $\chi^2(1)$ -distribution. The critical significance level is 83% which leads to the unambiguous test result that the null can not be rejected. The conclusion is strong support for the long-run inflation hedge hypothesis.

< Figures 3 and 4 >

To check the robustness of this conclusion, we have examined whether results are stable over time by estimating the cointegrated VAR model recursively. Figures 3 and 4 provide the results, showing the recursive estimates of the three eigenvalues and of the coefficients of the (one) cointegrating vector, respectively. The eigenvalues are fairly stable over the sample period, so the conclusion of one and only one cointegrating relation in the data is robust over time. Figure 4 shows that the long-run coefficients are reasonably stable, maybe with the exception of a slight instability of the coefficient to the discount rate in the late part of the sample. Most importantly, the coefficient to the price level is very stable. We take these results as evidence that the conclusion in favor of the inflation hedge hypothesis is robust over time.

The cointegrated VAR model approach has the advantages, compared to single-equation-cointegration methods, that it allows for more than one cointegrating relation in the data and, in general, leads to consistent and asymptotically efficient estimates of the long-run parameters (ϑ^*). However, as noted by e.g. Gonzalo and Lee (1998), Johansen (1999) and Juselius (1999), the cointegrated VAR model is sensitive to the number of observations. Thus, evidence based on Monte Carlo simulations suggests that the test of cointegration and the tests of hypotheses on the long-run coefficients may suffer from poor small sample performance (size distortions and low power). Moreover, inference from the model is based on the condition that the VAR specification gives the correct model not only for the variable of interest (stock prices) but also for the remaining variables (the general price level and the discount rate). As a further check on the robustness of conclusions, we have therefore re-estimated (5) by single-equation-cointegration methods. These give valid and efficient

inference in our case because we only have one cointegrating relation and because there is only error-correction in the direction of stock prices, implying that the price level and the discount rate are weakly exogenous for the parameters of the cointegrating vector, cf. Johansen (1996, Chp. 8). The latter can be shown by formal testing¹⁴.

Given the evidence of cointegration, OLS estimation of (5) produces consistent estimates of the coefficients¹⁵. However, testing coefficient hypotheses based on these estimates is in general difficult due to a (possible) correlation between the error term in the cointegrating relation and the innovations in the regressors, cf. Hamilton (1994). In particular, usual t-test statistics calculated from the OLS coefficients and the OLS standard errors do not have standard (known) distributions. Therefore, we have to refine the estimation of the cointegrating relation. Several approaches have been suggested for this purpose, cf. e.g. Phillips and Loretan (1991), Stock and Watson (1993) and Phillips and Hansen (1990). Hamilton (1994) and Mills (1993) provide surveys. We employ two of these procedures, both suggested by Phillips and Loretan (1991); the Phillips-Loretan OLS procedure (PLOLS) and the Phillips-Loretan Non-linear least squares procedure (PLNLS).

In both approaches, the static regression of (5) is augmented by stationary terms which capture the short-run dynamics of the explanatory variables. The PLOLS procedure augments (5) with current, lagged and leaded first differences of the explanatory variables (the price level and the discount rate), leading to the dynamic regression:

$$(10) \quad q_t = \mathbf{b}_0 + \mathbf{b}_1 p_t + \mathbf{b}_2 t + \mathbf{b}_3 r_t + \sum_{i=-N_1}^{N_1} \mathbf{g}_{1i} \Delta p_{t-i} + \sum_{i=-N_2}^{N_2} \mathbf{g}_{2i} \Delta r_{t-i} + u_t$$

t denotes as before the deterministic time trend (replacing y_t in (5)) and u_t is the new residual term. N_1 and N_2 which determine the number of lags and leads in the regression have to be

¹⁴ We have weak exogeneity if the equilibrium error does not affect the short-run dynamics of the price level and the discount rate, i.e., if the corresponding adjustment coefficients in α (see first column, second and third entry of α in Table 3) are both zero. This hypothesis can be tested formally by a LR test. The LR test statistic is 1.04 which has to be compared with a $\chi^2(2)$ -distribution. The critical significance level is 59%, leading to the conclusion that weak exogeneity can not be rejected.

specified prior to estimation. We use different specifications, cf. below, in order to check the sensitivity of coefficient estimates. (10) is estimated by OLS.

The PLNLS procedure augments (5) further by adding lagged levels of the error correction term, i.e., the difference between stock prices and their long-run equilibrium level as determined by (5), $[q_t - (b_0 + b_1 p_t + b_2 t + b_3 r_t)]$:

$$(11) \quad q_t = b_0 + b_1 p_t + b_2 t + b_3 r_t + \sum_{i=-N_1}^{N_1} g_{1i} \Delta p_{t-i} + \sum_{i=-N_2}^{N_2} g_{2i} \Delta r_{t-i} + \sum_{i=1}^{N_3} f_i (q_{t-i} - b_0 - b_1 p_{t-i} - b_2(t-i) - b_3 r_{t-i}) + v_t$$

The error correction terms are included in order to eliminate serial correlation in the disturbance term (v_t) and increase the efficiency of the coefficient estimates, cf. Hamilton (1994). Because the coefficients of the cointegrating relation enter the lagged error correction terms, (11) is estimated by Non-linear least squares (NLS).

< Table 4 >

Results including t-tests on the price level coefficient are reported in Table 4. In the first entry, results from estimating (5) by OLS (no augmentation) are shown together with OLS standard errors which are indicative only. The PLOLS procedure is used in three regression specifications which differ according to the included first differences of the price level and the discount rate (entries 2 through 4). In the first application (entry 2), current first differences and first differences at lead 1 and lag 1, respectively, are included. The disturbance term shows serial correlation up to lag 5 so standard errors and t-statistics have to be corrected. We use the adjustment method suggested by Hamilton (1994), based on an AR(5)-model fitted to the residuals of the PLOLS regression¹⁵. In the second application (entry 3), first differences of up to 2 leads and 2 lags are included. This further augmentation only has a minor effect on the estimated price level coefficient. It turns out that the disturbance term

¹⁵ Using the two-step procedure of Engle and Granger (1987), we can, at the 10% significance level, confirm (5) as a cointegrating relation, see Appendix A (the alternative based on CPI and a trend).

¹⁶ The adjusted t-statistics reported in Table 4 (entry 1) are calculated as the ordinary OLS t-statistics multiplied by the ratio (s/l), where s is the ordinary standard error of the residual in (10) while l is calculated from an AR(5)-model fitted to the residual, see Hamilton (1994, p. 610). l can, heuristically, be interpreted as an estimate of the residual standard error in 'long-run equilibrium' of the AR(5)-model. The reported standard errors of the coefficient estimates are adjusted accordingly.

shows no sign of misspecification in this formulation (no serial correlation) so usual OLS standard errors can be used. Finally, in the third application (entry 4), we use a “specific-to-general” procedure and augment (5) with current, lagged and leaded first differences until the disturbance term fulfills the white noise requirements. The resulting regression is just a reduced version of the second PLOLS regression (entry 3) where insignificant first difference terms have been omitted. Again, the estimated price level coefficient is only mildly affected. PLOLS regressions have also been carried out with more leads and lags but the coefficients and, in particular, the price level coefficient are stable wrt. this further augmentation.

The PLNLS regression in entry 5 has the augmenting terms shown in the first column of the table, including one lag of the error correction term. The augmenting terms are chosen in a “specific-to-general” manner in order to ensure a white noise disturbance. The reported standard errors are NLS calculated standard errors.

The results show that while the coefficient estimates for the trend and especially the discount rate are sensitive to the estimation procedure used, the estimate of the price level coefficient is fairly robust (and also comes close to the estimate obtained from the cointegrated VAR model). Turning to the inflation hedge hypothesis, the t-tests show that the price level coefficient is significant in all four cases. Moreover, in none of the cases we can reject the hypothesis that the price level coefficient is 1. The evidence in terms of critical significance levels is very strong. Hence, we conclude that single-equation-cointegration methods confirm the strong evidence in favor of the hedge hypothesis.

5. Conclusion and Discussion

We have examined whether Danish stocks provide a hedge against inflation, focusing explicitly on the long-run horizon. We have tested the hypothesis based on the long-run relation between stock prices and the general price level, estimated by cointegration analysis. Using the Consumer Price Index as the relevant price measure, results give strong support to the hedge hypothesis. The evidence supports the hedge property in its most restrictive sense of a perfect hedge. The conclusion is confirmed by both multivariate and univariate cointegration methods and is robust over time. The inflation hedge hypothesis is tested within

a firm modeling framework which is validated by the data as a cointegrating relation for stock prices.

The inflation hedge property of stocks (defined as a perfect hedge) only receives weak support, if any, in the literature. The strong support in this paper is therefore not a standard result. We do not believe that the Danish stock market has unique characteristics compared to other stock markets but rather attribute the difference to the literature to other factors. First of all, the use of different investment horizons is one possible explanation. We test the inflation hedge hypothesis in a long-run framework whereas others, e.g. Fama and Schwert (1977) and Gultekin (1983), examine relatively short investment horizons (less than 6 months). A plausible and reconciling interpretation of this evidence is that stocks hedge against inflation in the long run, but not in the short run.

Second, the use of different sample periods may be important. In this paper, we include observations until 1996, while other studies, e.g. Fama (1981) and Gultekin (1983), use samples that only cover the period until the end of the 1970s. As well-known, the 1970s were in almost all OECD countries a period of very high and increasing inflation due to the 1973 and 1979 oil price shocks. The use of a sample ending shortly after the oil price shocks ignores the subsequent and major adjustment in stock prices and may have triggered the (false) conclusion that stocks do not hedge against inflation. In this context, it may in particular be important that real oil prices, while increasing substantially during the oil crises with a deteriorating effect on profit margins, have by the beginning of the 1990s returned to the pre-oil crises level, hence allowing for a restoration of “normal” profit margins. Our study differs from the older literature by including the important adjustment period after the 1970s.

Finally, we have taken a different approach compared to the literature where it has been standard to test the inflation hedge hypothesis based on return regressions. We use cointegration methods to disentangle the short-run dynamics of and the long-run linkages between stock prices and the general price level, explicitly allowing for slow adjustment in stock prices to long-run equilibrium in the event of shocks to (not least) the general price level. This approach has the advantage of allowing for a clear identification of long-run stock

price behavior. The return regressions approach, on the other hand, does not distinguish between short-run dynamics and long-run linkages and the identification of long-run stock price behavior is conducted merely by investigating a sufficiently ‘long’ investment horizon. However, this muddles short-run dynamics and long-run linkages. Moreover, by linking stock returns to contemporaneous inflation, return regressions from the outset preclude slow adjustment in stock prices. In principle, this could trigger a false conclusion that stocks do not hedge against inflation in the long run. That is, stocks may be a perfect hedge against inflation with a lagged response in stock prices, but return regressions may fail to establish this because they do not explicitly take account of the lagged adjustment. As an illustration, assume that stocks hedge against inflation after a lagged adjustment over (say) 3 years, i.e., stock prices adjust completely to current inflation after 3 years. A return regression for even a long investment horizon of e.g. 5 years may not be able to detect this because stock returns do not reflect (completely) inflation in the last 3 years of each horizon, while at the same time, stock returns in the first 3 years are a result of adjustment to inflation in the preceding years¹⁷.

To highlight the difference between the standard return regressions approach and our approach (the cointegration approach) more formally, consider the cointegrated VAR model of (8) with a lag length of (for simplicity) $k=1$ and let us assume that this is the ‘true’ reduced form model. The structural form of the cointegrated VAR model is formally derived by premultiplying this reduced form by a non-singular matrix, cf. Johansen (1996). The resulting dynamic equation for stock prices can be written as (ignoring the disturbance term):

$$(12) \quad \Delta q_t = a_0 + a_1 \Delta p_t + a_2 \Delta r_t + a_3 \beta^{*'} X_{t-1}^*$$

where the a_i ’s are structural coefficients. The term $\beta^{*'} X_{t-1}^*$ denotes as before the error-correction term from the long-run stock price relation (as of period $t-1$). Now notice, that the

¹⁷ The possibility of a slow adjustment in stock prices or rather stock returns to a change in the inflation rate has also been noted by Barnes *et al.* (1999). They test the inflation hedge hypothesis on a large sample of countries using the standard return regressions approach. To take account of the possible slow adjustment, they include both the contemporaneous and the lagged inflation rate in the return regressions. However, this does not alter the evidence significantly. The general result in Barnes *et al.* (1999) is a rejection of the inflation hedge hypothesis for stocks.

first differences part of (12) resembles a return regression for an investment horizon of 1 year, by regressing the first differences of log-to-stock prices (Dq_t), which is a proxy for the 1-year stock return, on 1-year inflation (Dp_t) and the 1-year change in the discount rate (Dr_t). The 1-year return regression, therefore, can be viewed as a special case of (12) where the level term (the cointegration term) has been excluded. In terms of (12), what distinguishes the cointegration approach from the standard approach is that the former is concerned with the long-run coefficients to stock prices and the general price level, i.e., the cointegrating vector ϑ^* . The return regressions approach, on the other hand, is concerned with the dynamic coefficient to the price level, i.e., the coefficient a_1 which captures the short-run or contemporaneous response in stock prices to inflation. This difference reflects our explicit focus on the long-run horizon whereas the existing literature has mainly examined the inflation hedge hypothesis over relatively short horizons.

(12) also suggests a possible shortcoming of the standard approach. In standard return regressions, the level term of (12) is excluded which (implicitly) assumes either that stock prices adjust immediately to their long-run equilibrium level as determined by the cointegrating relation (i.e., the equilibrium error $\vartheta^* X_{t-1}^*$ is always zero), or that there is no cointegration between the level variables (i.e., the cointegrating rank is zero and no cointegrating vectors ϑ^* exist). In our case, both possibilities are rejected by the data. Therefore, the 1-year return regression must be misspecified because it omits a significant regressor (the level term of (12)), which reflects the slow adjustment in stock prices. In general, the result is inconsistent coefficient estimates, which affects the inference on the coefficient to inflation and, hence, the inflation hedge hypothesis. In the cointegration approach, we explicitly allow for slow adjustment¹⁸.

For the purpose of comparison, we have also tested the inflation hedge hypothesis for Danish stocks using standard return regressions over the same sample period as considered above.

We have run a return regression where nominal stock returns are regressed (OLS) on

¹⁸ The cointegrated VAR model is more general than implied by (12) because it allows for more short-run dynamic terms, i.e., lagged first differences, when the lag length k is larger than 1. Moreover, in the case of a cointegrated VAR model with an explicit measure for real production, we would also have included the 1-year real growth rate (Dy_t) as a first-differences regressor in (12). Note that (12) focuses on the 1-year investment horizon. For longer horizons, the implied model for returns will be more complicated but the fundamental insight remains that return regressions omit significant level terms.

contemporaneous inflation and a constant term for each of the five investment horizons of 1 to 5 years duration¹⁹. The point estimates of the coefficient to inflation range from 0.27 (1-year horizon) to 0.94 (5-year) so the estimates for the longest horizons come close to one, the value consistent with the inflation hedge hypothesis. We can test whether the response in stock returns to inflation is statistically significant. Using the Newey and West (1987) coefficient standard error, which is consistent to heteroskedasticity and serial correlation in the regression residual (up to lag 5), we get a t-test statistic of 1.5 for the null hypothesis that the coefficient to inflation is insignificant (zero) at the 5-year horizon (where the inflation coefficient is most significant). The conclusion is that, even though the point estimate is high and close to one, the estimation uncertainty is substantial and the inflation effect is, in statistical terms, only weakly significant. Using conventional significance levels, we would accept the null that inflation has no effect on stock returns. At best, these results only provide weak support to the inflation hedge hypothesis. Moreover, it turns out that the high point estimates of the inflation coefficient hinge primarily on two events, that is, the exceptionally large stock returns encountered in the years 1972 and 1983 (see Figure 1 for the large capital gains on stocks in these two years). If we eliminate the influence of the returns realized in these two years by including dummies in the regressions of stock returns on contemporaneous inflation, the resulting point estimates of the inflation coefficient are substantially lower and now range from -0.11 (1-year horizon) to 0.21 (2-year)²⁰. The coefficient at the 5-year horizon is estimated to be 0.14. None of these coefficients are significantly different from zero.

¹⁹ The stock return is annualized, discretely compounded and includes capital gain and dividend yield, cf. Nielsen, Olesen and Risager (1997). Inflation is measured by the continuously compounded annual growth rate in the Consumer Price Index (CPI). The return regressions use overlapping observations for each of the horizons of 2 to 5 years. Appendix B provides further evidence from using the return regressions approach to test the hedge hypothesis, including results for the 10-year horizon and results from an extended return regression with additional explanatory variables. At first sight, the results in Appendix B give more support to the hedge hypothesis than the regression results presented here. In particular, the results for the 10-year horizon seem to confirm the hedge hypothesis. However, a closer examination raises serious doubts about the reliability of this conclusion, in particular, because the return regressions suffer from highly unstable coefficient estimates over the sample, including an unstable estimate of the coefficient to inflation.

²⁰ The regressions include two impulse dummies that eliminate *all* the effects of the stock returns in 1972 and 1983. That is, for the 2-year regression where we use overlapping observations, we include one dummy which has the value of one for 1972 and 1973 and is zero otherwise, and another dummy which has the value of one for 1983 and 1984 and is zero otherwise. Similarly, for the 3-year regression each dummy takes on the value of one three years in a row, and so forth. We want to exclude the returns of these two years because they represent clear outliers (annual returns of 95% and 118%, respectively) which can, furthermore, be explained

Hence, using the cointegration approach gives stronger and more reliable support for the inflation hedge hypothesis than produced by standard methods. This evidence suggests that the differences in approach could be important in understanding why we find much stronger support for the hedge hypothesis than in the literature. On this background, and recalling the possible shortcomings of the return regressions approach, cf. above, an obvious topic for future research would be to use the cointegration approach (and a more recent sample) to re-examine the evidence on the inflation hedge hypothesis for other stock markets.

In testing the inflation hedge hypothesis, we have focused exclusively on stock prices, thereby ignoring dividends as part of overall stock returns. In principle, a hedge against inflation demands that the total stock portfolio consisting both of stocks bought at the time of initial investment and stocks bought subsequently by the reinvestment of dividends retains a stable purchasing power in the event of inflation. Our stock price approach takes a shortcut by focusing on the value of initial stocks only. One reason is that we want to test the hedge hypothesis within a firm modeling framework which in our case is a model for stock prices. Moreover, it can be argued that for all practical purposes the exclusion of dividends is not important because dividend yields are fairly modest (in particular in recent history). Thus, whether the neglect of dividend payments is of importance is reflected by how much future reinvestments amount to as a fraction of total future portfolio value. This, again, is determined by the dividend yield (in absolute terms). Because the dividend yields for Danish stocks over the sample period considered are small (between 1% and 7%), the stock price approach should give a fairly good approximation to the ‘overall’ inflation hedge question in terms of total portfolio value.

We have, as standard in the literature, ignored costs of stock transactions and investor taxes on stock returns, i.e., dividend and capital gains taxes. While the neglect of transaction costs may not be so important because the hedge property of stocks relate to a passively held portfolio with limited active trading, it is a more open question whether the neglect of taxes matters. Dividend taxes should not matter because the behavior of stock prices is the crucial

by one-off exceptional changes in the Danish economy such as the Danish referendum in 1972 leading to

issue for the inflation hedge property, cf. above. In the case of a capital gains tax it could, tentatively, be argued that the inflation hedge property of stocks is retained on an after-tax basis as long as the tax is fixed and proportional because the ‘after-tax’ stock price (or portfolio value) would be proportional to the ‘before-tax’ stock price, thereby preserving a one-for-one relation with the general price level. However, capital gains taxes are not proportional and have not been fixed over time, so this aspect needs closer investigation. This is an interesting but in the Danish case also highly challenging question to address because taxation rules are complex and differ markedly between different types of investors.

The result that (Danish) stocks in the long run hedge against inflation should be of interest to any rational investor who cares about the real value of his investments and who has a long-run investment horizon, e.g. a pension saver. Hence, the hedge property has *ceteris paribus* implications for optimal portfolio choice because not all assets are immune to inflation uncertainty. For instance, nominal bonds can at most compensate for expected inflation, leading to real uncertainty of a bond investment. However, one should be careful with the proper interpretation of the inflation hedge result.

First of all, the result relates to the market portfolio of stocks, i.e. the highly-diversified portfolio consisting of all stocks listed at the CSE. A high degree of diversification must be expected to be necessary in order to sustain the hedge property against the overall price level because relative prices of goods and services and, hence, relative firm profits change over time.

Secondly, the inflation hedge is a long-run phenomenon, so that stocks should be expected to compensate for inflation in the ‘long run’ and the ‘long run’ only. Our analysis does not answer the question how long the ‘long run’ is, but the lag length of the VAR model and the estimated adjustment parameter (together with return regressions, see above and Appendix B) loosely indicate that a time span of 5-10 years is the appropriate horizon. It should be emphasized that the interpretation is not that stocks compensate for contemporaneous inflation over a fixed (say) 5-year horizon, but rather that stock prices have adjusted

membership of the EEC and the introduction of a new, separate pension fund tax on bond returns in 1983.

completely to current inflation after an adjustment period of 5 years. Thus, over a fixed investment horizon of 5 years an investor may only get compensated for the first year of inflation but not (fully) for the remaining 4 years of inflation. This interpretation distinguishes our approach from the return regressions approach which focuses on a fixed horizon, and is related to the different assumptions on the adjustment-speed of stock prices.

Finally, the hedge result is framed within a structural model for stock prices where real production and a discount rate also enter as explanatory variables. The result that stock prices move one-for-one with the general price level therefore applies to what could be called a controlled or *ceteris paribus* ‘experiment’ where the price level is changed while real production and the discount rate are kept fixed. This means that in the event of inflation, stock prices may not always end up by increasing proportionately, i.e., with the same relative change as the general price level, because (and only because) real production and the discount rate may have changed simultaneously. This does not contradict the hedge hypothesis but rather reflects the fact that stock prices do not depend only on the price level. Thus, real stock prices may change due to innovations in real fundamentals. The standard return regression literature also focuses on a *ceteris paribus* ‘experiment’ when drawing inference on the inflation hedge hypothesis.

In the field of monetary economics, it is standard to distinguish between the ‘neutrality’ and ‘superneutrality’ of money, see e.g. Grandmont (1988). ‘Neutrality’ means that a change in the *level* of the money stock has no effects on real economic variables (production, employment, and so forth) whereas ‘superneutrality’ implies that a change in the *growth rate* of the money stock has no real consequences. Using a corresponding terminology, the evidence provided on the long-run relationship between stock prices and the general price level can be given the interpretation that inflation is both neutral and superneutral to stocks in the long run. That is, a permanent change in the general price *level* will, according to (9), eventually lead to a proportionate change in stock prices, leaving real stock prices unaffected. Similarly, from (9), a permanent change in the *rate of inflation*, i.e., the growth rate in the general price level, will in the long run result in an equivalent change in the growth rate of stock prices and there will be no impact on real stock prices. Notice that it is a prerequisite

for both neutrality and superneutrality that real production and the real discount rate are unaffected. This seems as the reasonable assumption when considering the long-run responses to one-off changes in the price *level*. As to the case of superneutrality, changes in the inflation rate will have no impact on real production and the discount rate and, hence, real stock prices *if* the economy is characterized by the property of Classical Dichotomy between the real and the money sectors in the long run²¹. The latter assumption is often used both in macroeconomic theory (it forms, in particular, the corner stone of the traditional Neoclassical-Keynesian Synthesis) and applied business-cycle research.

²¹ Following the definition by Grandmont (1988, p.2), Classical Dichotomy applies if real magnitudes are determined exclusively by the real sector while absolute prices are determined by the equilibrium condition for money.

References

- Barnes, M., J.H. Boyd and B.D. Smith (1999), "Inflation and Asset Returns", *European Economic Review* 43, 737-754.
- Bonnichsen, R. (1983), "Aktiemarkedet og købekraften: Kantotalafkastet af aktier følger med inflationen?", *FINANS/INVEST* 82 (8), 58-63 (in Danish).
- Boudoukh, J. and M. Richardson (1993), "Stock Returns and Inflation: A Long-Horizon Perspective", *American Economic Review* 83 (5), 1346-1355.
- Campbell, J.Y., A.W. Lo and A.C. MacKinlay (1997), *The Econometrics of Financial Markets*, Princeton University Press.
- Dickey, D.A. and W.A. Fuller (1979), "Distribution of the Estimators for Autoregressive Time Series with a Unit Root", *Journal of the American Statistical Association* 74, 427-431
- Doornik, J.A. and H. Hansen (1994), "A Practical Test of Multivariate Normality", Working Paper, Nuffield College.
- Doornik, J.A. and D.F. Hendry (1997), *Modelling Dynamic Systems Using PcFiml 9.0 for Windows*, International Thomson Business Press (Manual for PcFiml).
- Ely, D.P. and K.J. Robinson (1997), "Are Stocks a Hedge Against Inflation ? International Evidence Using a Long-Run Approach", *Journal of International Money and Finance* 16 (1), 141-167.
- Engle, R.F. and C.W.J. Granger (1987), "Co-integration and Error Correction: Representation, Estimation and Testing", *Econometrica* 55 (2), 251-276.

- Engle, R.F. and C.W.J. Granger (editors) (1991), *Long-Run Economic Relationships: Readings in Cointegration*, Oxford University Press.
- Fama, E.F. (1981), “Stock Returns, Real Activity, Inflation and Money”, *American Economic Review* 71 (4), 545-565.
- Fama, E.F. and G.W. Schwert (1977), “Asset Returns and Inflation”, *Journal of Financial Economics* 5, 115-146.
- Faust, J. and E.M. Leeper (1997), “When Do Long-Run Identifying Restrictions Give Reliable Results?”, *Journal of Business and Economic Statistics* 15 (3), 345-353.
- Frennberg, P. and B. Hansson (1993), “Stock Returns and Inflation in Sweden. Do Stocks Hedge against Inflation”, Working Paper #25, Department of Economics, Lund University.
- Gonzalo, J. and T.H. Lee (1998), “Pitfalls in Testing for Long Run Relationships”, *Journal of Econometrics* 86, 129-154.
- Grandmont, J.-M. (1988), *Money and Value*, Cambridge University Press, New York.
- Gultekin, N.B. (1983), “Stock Market Returns and Inflation: Evidence from Other Countries”, *Journal of Finance* XXXVIII (1), 49-65.
- Hamilton, J.D. (1994), *Time Series Analysis*, Princeton University Press.
- Hansen, H. and S. Johansen (1996), “Recursive Analysis of Eigenvalues in Cointegrated VAR-Models”, revised version of Discussion Paper 92-13, Institute of Economics, University of Copenhagen.
- Johansen, S. (1996), *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*, Oxford University Press.

Johansen, S. (1999), “A Bartlett Correction Factor for Tests on the Cointegrating Relations”, European University Institute Working Paper ECO 99/10.

Juselius, K. (1999), “Models and Relations in Economics and Econometrics”, Working Paper, University of Copenhagen.

Kaul, G. (1987), “Stock Returns and Inflation: The Role of the Monetary Sector”, *Journal of Financial Economics* 18, 253-276.

Kwiatkowski, D., P.C.B. Phillips, P. Schmidt and Y. Shin (1992), “Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root”, *Journal of Econometrics* 54, 159-178.

MacKinnon, J.G. (1991), “Critical Values for Cointegration Tests”, Chp. 13 in Engle and Granger (eds.) (1991).

Mills, T.C. (1993), *The Econometric Modelling of Financial Time Series*, Cambridge University Press.

Newey, W.K. and K.D. West (1987), “A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix”, *Econometrica* 55, 703-708.

Nielsen, S., J.O. Olesen and O. Risager (1997), Danish Stock Market and Macroeconomic Database, Department of Economics, Copenhagen Business School.

Phillips, P.C.B. and B.E. Hansen (1990), “Statistical Inference in Instrumental Variables Regression with I(1) Processes”, *Review of Economic Studies* 57, 99-125.

Phillips, P.C.B. and M. Loretan (1991), “Estimating Long-Run Economic Equilibria”, *Review of Economic Studies* 58, 407-436.

Phillips, P.C.B. and P. Perron (1988), “Testing for a Unit Root in Time Series Regression”, *Biometrika* 75, 335-346.

Reimers, H.E. (1992), “Comparisons of Tests For Multivariate Cointegration”, *Statistical Papers* 33, 335-359.

Sellin, P. (1998), “Monetary Policy and the Stock Market: Theory and Empirical Evidence”, Working Paper #72, Sveriges Riksbank.

Stock, J.H. and M.W. Watson (1993), “A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems”, *Econometrica* 61, 783-820.

Appendix A: Alternative Candidates for a Cointegrating Relation

This appendix reports the results from estimating alternative candidates for a cointegrating relation for stock prices, cf. (5), while using alternative measures of the general price level and real production, respectively. For the general price level, we have examined three different measures, the official Consumer Price Index (CPI) (denoted by p_t in the following), the implicit price deflator for total GDP in factor prices (pyf_t) and, finally, the implicit price deflator for GDP in factor prices in the sector of manufacturing (pyf_{it}). For real production, we use data on total GDP (denoted by yf_t), respectively, GDP in manufacturing (yf_{it}), both in fixed 1980-factor prices. The price deflators and production measures are taken from National Accounts and all series are in log-levels. For the growth-adjusted real discount rate (r_t in (5)), we, throughout, use the same proxy as in the main text, cf. section 3. The estimations and tests are performed by a single-equation cointegration method, that is, the Engle and Granger (1987) two-step procedure (EG2).

To begin with, we have to test for the stationarity properties of the data series (stock prices, all price and production measures and the discount rate proxy). Unit root tests have been performed following the same approach as in section 3 and the conclusion is that all series are integrated of order 1, i.e., non-stationary in levels but stationary in first differences (tests not reported). Hence, the regression in (5) is balanced which is a prerequisite for using the EG2 procedure for estimation purposes.

< Table A.1 >

Table A.1 reports the alternative estimates of (5) and the corresponding tests for cointegration. The measures used for the general price level and real production are indicated in the first column. For example, the regression of the first entry uses the price deflator for total GDP and, correspondingly, total GDP in fixed prices as the relevant measures. The second column shows the OLS estimates of the coefficients of (5) (stated as a cointegrating vector which is normalized on stock prices), together with *indicative* OLS standard errors. For instance, the price level coefficient in the regression of the first entry (the estimated ϑ_1) is 1.38 with an indicative OLS standard error of 0.14. A residual-based test for cointegration is

performed by testing the null hypothesis of a unit root in the process for the OLS residuals. If the null is rejected, the residuals are stationary and the regression (5) is concluded to be a cointegrating relation. The test results and conclusions on cointegration are reported in the remaining columns of the table, using the cointegrating regression Dickey-Fuller test (CRDF), cf. Engle and Granger (1987) or Hamilton (1994). The number of augmenting lags (of the first differences of the residuals) in the CRDF test is chosen according to a “specific-to-general” procedure, taking the simple Dickey-Fuller regression without augmentation as the starting point and - in case this regression shows signs of being misspecified (serial correlation in the disturbance term) - including lags until diagnostic tests are passed. A maximum of 1 augmenting lag suffices in the tests reported in Table A.1²².

The choice of production measure is important for whether or not (5) is a cointegrating relation. The first two regressions in Table A.1 which both use an explicit measure of production show no cointegration at the 10% significance level. Because the OLS estimates *indicate* that the production measures are insignificant, a regression is run (third entry) where prices and production are combined in nominal production (using total GDP in current prices as the relevant measure, denoted by y_t^* in the table) to check whether this enhances the presence of cointegration. This is not the case²³.

In the last three regressions in Table A.1, we have replaced the explicit production measure by a deterministic trend (denoted by t), which can be interpreted as a proxy for the trend growth in production. Results show that the inclusion of a deterministic trend leads to cointegration at the 10% significance level when measuring the general price level by CPI (entry 4) or the factor price deflator for manufacturing (entry 6). Cointegration is most evident in the latter case with cointegration being accepted also at the 5% significance level.

²² The level of augmentation used in Table A.1 is the same as one would get from a “general-to-specific” procedure, starting out with a Dickey-Fuller regression with 5 augmenting lags and then, successively, removing insignificant lags from the highest order.

²³ We have examined alternative measures of real production including GDP for the private sector and GDP for the private sector excluding farming and housing, but without any further success. The lack of cointegration and the apparent insignificance of the production measures, basically, suggests that we have not been able to find a good proxy for the production of goods and services by the representative firm on the Copenhagen Stock Exchange.

Cointegration is just rejected when using the factor price deflator for total GDP as the price measure (entry 5).

Despite the fact that the evidence of cointegration is strongest when using the price deflator for manufacturing as the price measure, we prefer to test the hedge hypothesis in terms of CPI inflation for the reasons stated in section 3²⁴.

²⁴ While CPI seems most relevant to stock investors, a deflator for GDP in factor prices (or a net price index) may actually be more relevant to firm profits and, hence, more adequate for the theoretical framework of section 2. Thus, CPI includes indirect taxes paid by the consumers. Furthermore, CPI measures the prices of (domestically consumed) consumer goods and services and, thereby, ignores (say) the prices of investment goods. A factor price deflator captures the prices of all goods and services produced. However, whatever price measure used, it is just a proxy for what we really want to measure, and that is the prices of goods and services produced by the representative firm at the Copenhagen Stock Exchange. In particular, what we need is a good proxy for the long-run movements in the 'true' prices and, in this respect, CPI may do as well as e.g. a factor price deflator. It should also be recalled that the use of a proxy for the price level does not undermine the asymptotic consistency of the coefficient estimates in a cointegrating relation, cf. Hamilton (1994), provided cointegration is preserved.

Appendix B: Evidence from Return Regressions

In the literature, the inflation hedge hypothesis is tested by examining the link between stock returns and contemporaneous inflation. This appendix provides comparable results for Denmark, focusing on the three investment horizons of 1, 5 and 10 years.

The Empirical Model

Based on the theoretical framework of section 2, we use the following empirical model for stock returns over the k -year investment horizon²⁵:

$$(B1) \quad Sk_t = b_0 + b_1 Ik_t + b_2 GYk_t + b_3 GRk_t + e_t, \quad k = 1, 5 \text{ and } 10 \text{ years}$$

Sk_t denotes the annualized total stock return over the k -year investment horizon, including both capital gains and dividend yield. $Ik_t \equiv (\Delta_k \ln P_t) / k$ and $GYk_t \equiv (\Delta_k \ln Y_t) / k$ are, respectively, the (continuously compounded) annual inflation rate and the (continuously compounded) annual growth rate in real production over the same k year horizon.

$GRk_t \equiv (\Delta_k \ln R_t) / k$ is the per annum relative change in the discount rate over the investment horizon while e_t , finally, denotes the usual disturbance term. According to (B1), stock returns should be regressed on a constant term and contemporaneous values of the inflation rate, the real growth rate and the relative change in the discount rate. Whether or not stocks provide a hedge against inflation is captured by the coefficient to inflation, β_1 , measuring the direct or partial effect from inflation to stock returns. A formal test of the hedge hypothesis is performed in two steps, by testing (i) whether inflation has a significant effect on stock returns ($\beta_1 \neq 0$), and, if the inflation effect is significant, (ii) whether the relationship between (changes in) stock returns and inflation, furthermore, is one-to-one ($\beta_1=1$).

The Data

²⁵ The theoretical counterpart of (B1) is obtained by taking k -year differences in the logarithmic analog to (3) and dividing through by k to obtain *per annum* continuous growth rates. We substitute total stock returns (including dividend yields) for capital gains as the endogenous variable to allow for a comparison with the literature. It can be shown that this does not affect the empirical results significantly. The latter reflects the fact that the variation in dividend yields have played only a minor role for the *variation* in Danish stock returns over the sample period.

Data for total stock returns are from the database by Nielsen, Olesen and Risager (1997) and relate to the market portfolio of all Danish stocks listed at the Copenhagen Stock Exchange. Inflation is measured by the (annualized and continuously compounded) growth in the official Consumer Price Index (CPI) by Statistics Denmark, while we as data for real growth use the (annualized and continuously compounded) growth in total GDP in fixed 1980 factor prices, taken from National Accounts. For the unobservable discount rate, we use the same proxy as in the main text, cf. section 3, and GRk_t is calculated as the (annualized) change in this proxy over a k -year period. All data are annual. In order to cover the same sample period as in the main text, we consider the period 1949-1996 for the 1-year investment horizon, 1953-1996 for the 5-year horizon and 1958-1996 for the 10-year horizon. At the 5- and 10-year horizons, we use overlapping observations. Figure B1 shows the stock return and the inflation at the three horizons.

< Figure B1 >

< Tables B1.a and B1.b >

Tables B1.a and B1.b report the outcome of tests for unit roots, using the Phillips and Perron (1988) Z_t -test (PP) and the test by Kwiatkowski *et al.* (1992) (KPSS). For the 1-year horizon, we conclude that all variables are stationary. The two tests give conflicting results for the real growth rate but the evidence seems most robust across lag lengths for the PP-test which strongly points to stationarity. For the inflation rate, the PP-test consistently concludes stationarity at the 10% significance level, while the KPSS test at the same time gives firm evidence in favor of stationarity when allowing for serial correlation in the disturbance term (lag length l^3). At the 5-year horizon, stock returns and the change in the discount rate ($GR5_t$) are stationary. Results for inflation are ambiguous as the PP-test points to (at least one) unit root whereas the KPSS test supports stationarity. Real growth is non-stationary according to both tests. Finally, for the 10-year horizon, results for both stock returns and inflation are ambiguous. $GR10_t$ is stationary while real growth is non-stationary. To conclude, the static regression of (B1) which requires the data series to be stationary is valid for the 1-year horizon whereas conclusions are less clear for the 5- and 10-year horizons. At the latter horizons, the real growth should be excluded to allow for a valid regression and the

regression results should, in general, be interpreted with caution due to the possible non-stationary behavior of the data series.

The Results

< Table B2 >

Results are shown in Table B2. Regressions of the type (B1) are performed for each of the three investment horizons. Furthermore, for each investment horizon three distinct regressions are examined, cf. below. For all regressions, OLS is used for estimating the parameters, producing consistent estimates. For the 5- and 10-year horizons standard errors of the parameter estimates are estimated by the Newey and West (1987) method to take account of heteroskedasticity and serial correlation up to lag 5 in the disturbance term. The non-standard behavior of the disturbance term can be motivated by the use of overlapping observations. The truncation at lag 5 seems appropriate as the Newey and West (1987) standard error of the inflation coefficient becomes stable at this lag length. For the 1-year horizon, we include impulse dummies for 1972 and 1983 in order to exclude the exorbitant and exceptionally high stock returns these years (returns of 95% and 118%, respectively). These outliers can be explained by exceptional changes in the Danish economy including, in particular, the Danish favorable EEC referendum in 1972, the major shift towards a new economic policy regime in late 1982 and the introduction of a new pension fund tax on bonds in 1983. Having included these dummies, the regression residual fulfills the white noise requirements of being serially uncorrelated and homoskedastic and, hence, standard errors of the coefficients can be estimated by OLS at the 1-year horizon.

The first regression for each horizon (first entry) shows the results for “the simple model”, which is the specification that has been used most extensively in the literature. This formulation is a special case of (B1) where any effects from real production and the discount rate are ignored ($\beta_2 \equiv \beta_3 \equiv 0$) so that stock returns are explained by inflation only. The problem with this formulation is that it, according to (B1), ignores potentially relevant explanatory variables. As well known, the omission of relevant regressors leads to biased estimates for the

remaining regressors to the extent that the omitted and included regressors are correlated. Thus, in the simple model, the estimated coefficient to inflation could potentially be biased, as it may also capture relevant effects from real growth and a changing discount rate. Results for the simple model should, therefore, in general be interpreted with caution.

The second regression for each horizon (second entry) shows results for (B1) including both real growth and the change in the discount rate (“the extended model”). The latter enters significantly and with the correct (minus) sign for all three horizons whereas real growth has the wrong sign (minus) in each case. However, the effect from real growth is also insignificant at the 5% significance level. For this reason, we exclude it from the regression (which is also preferable from unit root considerations, cf. above) and arrive at the “reduced extended model” (third entry for each horizon) which can be interpreted as the parsimonious model formulation. Comparing the reduced extended model with the simple one, we find that the inclusion of the discount rate matters at both the 5- and 10-year horizons as it increases the point estimate of the inflation coefficient.

In testing the inflation hedge hypothesis, we focus on the “reduced extended model” which provides the best specification in terms of included regressors. The impact of inflation on the stock return is clearly insignificant at the 1-year horizon, where the estimated inflation coefficient for all practical purposes is zero. At the 5-year horizon, the coefficient of 1.01 is very close to one, but the coefficient standard error is large (0.55) so that the inflation effect is only at the boarder of being significant, using conventional significance levels (the critical significance level of a two-sided t-test for significance is 6.6%). *If* the inflation effect is accepted to be significant, the hypothesis that stock returns and inflation move one-for-one is clearly accepted. At the 10-year horizon, the inflation coefficient of 1.01 is strongly significant with a t-statistic of almost 5. Moreover, the hypothesis that the inflation effect is one ($\vartheta_1=1$) receives strong support.

To judge the robustness of the latter evidence, a recursive estimation of the reduced extended model at the 10-year horizon is performed, cf. Figure B2.

< Figure B2 >

The support of the inflation hedge hypothesis at the 10-year horizon is certainly not stable over time. It is only with the inclusion of the observations in the mid-1980s that the hypothesis can be supported. Using a sample from 1958 to the early 1980s, the inflation effect is largely insignificant. This seriously questions the robustness of the inflation hedge result at the 10-year horizon. A similar picture can be shown for the 5-year horizon.

Conclusion

Using the standard return regressions approach, we find that stocks are certainly no hedge against inflation at the short 1-year horizon. Apparently, the hedge hypothesis receives mild support at the medium 5-year horizon and strong support at the long 10-year horizon. However, a closer examination raises serious doubts about the validity of this conclusion because estimates of the parameters in the return regressions and, in particular, the estimated coefficients to inflation, are highly unstable over time. Moreover, the regressions at the 5- and 10-year horizons suffer from a potential problem with non-stationary data series²⁶. Hence, we conclude that the return regressions approach does not produce reliable support to the hedge hypothesis.

²⁶ Other econometric problems include a small number of non-overlapping observations and a possible measurement error bias in coefficient estimates due to the use of a proxy for the discount rate regressor.

Figure 1. The Data

Levels and first differences of (by row and from the top) stock prices, the general price level (CPI) and the real discount rate (proxy). Variables in logs.

Sample: 1948-1996.




Figure 2. Diagnostic Graphics for the VAR Model

Actual and fitted values (in levels), residuals, residual correlogram and residual density for the equation for (by row and from the top) stock prices, the general price level and the real discount rate.


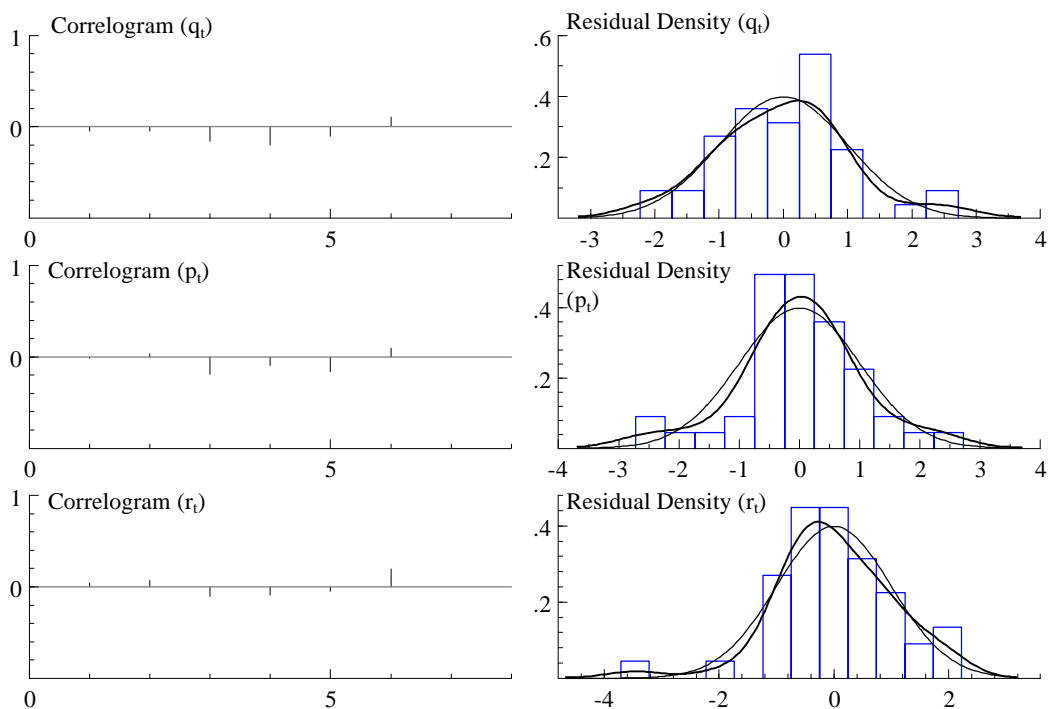


Figure 2, continued.



Note: Density plots include standard normal density for comparison (thin curve).

Figure 3. Recursive Estimates of the Eigenvalues of the Unrestricted VAR Model (Full sample: 1952-1996)

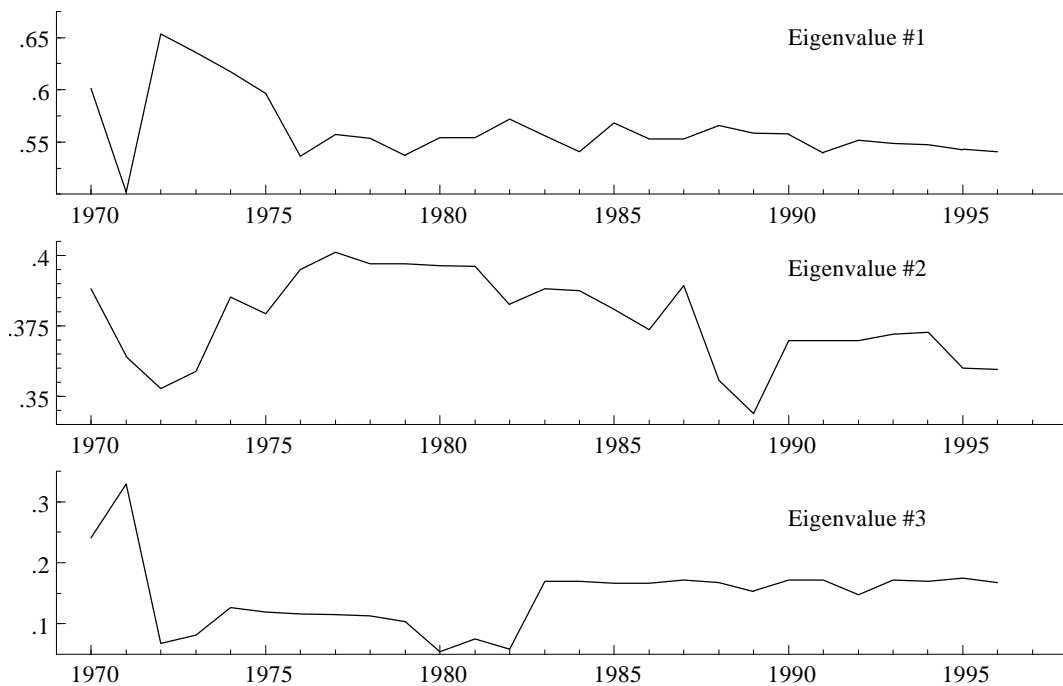



Figure 4. Recursive Estimates of the Cointegrating Vector

Recursive point estimates (solid line) and 95% confidence bands for the coefficient to (from the top) the general price level, the real discount rate and the deterministic trend. Cointegrating rank restricted to 1 and cointegrating vector normalized on stock prices. Full sample: 1952-1996.



Note (both figures): The method of recursive estimation keeps the short-run dynamics fixed at the one estimated for the entire sample, cf. Hansen and Johansen (1996) and Doornik and Hendry (1997).

Table 1.a. Phillips and Perron (1988) Z_t -Test for Unit Root
1948-1996

Series:	<i>Lag length (l)</i>						
	0	1	2	3	4	5	6
<u>Levels:</u>							
q_t (T)	-2.42	-2.24	-2.23	-2.24	-2.22	-2.23	-2.25
p_t (T)	-0.84	-1.08	-1.22	-1.34	-1.44	-1.53	-1.61
r_t	-2.68*	-2.64*	-2.61*	-2.59	-2.54	-2.53	-2.51
<u>First differences:</u>							
Δq_t	-8.95***	-8.97***	-9.04***	-9.13***	-9.29***	-9.43***	-9.52***
Δp_t	-2.79*	-2.94**	-2.80*	-2.70*	-2.73*	-2.78*	-2.86*
Δr_t	-8.14***	-8.15***	-8.25***	-8.41***	-8.69***	-9.00***	-9.45***
<u>Critical test values:</u>			10 %	5 %	2.5 %	1 %	
Without trend			-2.60	-2.93	-3.22	-3.58	
With trend			-3.18	-3.50	-3.80	-4.15	

Note: The Phillips and Perron (1988) unit root test is based on the first order autoregression $x_t = \alpha + \rho x_{t-1} + u_t$ (without trend), respectively, $x_t = \alpha + \rho x_{t-1} + \delta t + u_t$ (with trend) where the disturbance term u_t has mean zero but can otherwise be heterogeneously distributed (heteroskedastic) and serially correlated up to lag l , see also Hamilton (1994). The Z_t test statistic is as a modified t-statistic for the null hypothesis of a unit root ($\rho=1$), corrected for the possible non-standard properties of u_t . The null is rejected in favor of the stationary alternative ($\rho<1$) if Z_t is negative and sufficiently large in absolute value. Critical values are from Hamilton (1994, Table B.6) for a sample size of 50. *, ** and *** denote rejection of a unit root at the 10%, 5% and 1% significance level, respectively. All regressions include a constant term. (T) indicates that the regression includes a deterministic trend.

Table 1.b. Kwiatkowski *et al.* (1992) Test for Unit Root
1948-1996

Series:	<i>Lag length (l)</i>						
	0	1	2	3	4	5	6
<u>Levels:</u>							
q _t (T)	0.84***	0.48***	0.35***	0.28***	0.24***	0.21**	0.19**
p _t (T)	0.70***	0.36***	0.25***	0.19**	0.16**	0.14*	0.13*
r _t	2.58***	1.47***	1.07***	0.86***	0.73**	0.65**	0.58**
<u>First differences:</u>							
Δq _t	0.12	0.17	0.19	0.21	0.24	0.26	0.27
Δp _t	0.61**	0.36*	0.28	0.23	0.20	0.17	0.16
Δr _t	0.05	0.06	0.07	0.09	0.11	0.12	0.15
<u>Critical test values:</u>			10 %	5 %	1 %		
Without trend			0.35	0.46	0.74		
With trend			0.12	0.15	0.22		

Note: The Kwiatkowski *et al.* (1992) test for a unit root is a Lagrange Multiplier test of the null hypothesis that the series can be described by a stationary process (possibly around a deterministic trend), against the alternative that the process also includes a random walk component. The null of stationarity is rejected in favor of the unit root alternative if the test statistic is sufficiently large. Critical values are from Kwiatkowski *et al.* (1992). *, ** and *** denote rejection of the null (i.e., a unit root is accepted) at the 10%, 5% and 1% significance level, respectively. Lag length *l* is the number of lags allowed for in the stationary component of the process. (T) after a series indicates that the test allows for a deterministic trend, i.e., the null hypothesis is trend-stationarity. Otherwise, the null is mean-stationarity.

Table 2. Specification Tests of the VAR Model
Estimation sample 1952-1996

<u>Multivariate tests:</u>			
Vector Autocorrelation order 2	F(18,65) =	1.05	[0.42]
Vector Autocorrelation order 4	F(36,50) =	0.92	[0.60]
Vector Autocorrelation order 6	F(54,33) =	0.98	[0.54]
Vector Heteroskedasticity (squares)	F(156,2) =	0.02	[1.00]
Normality	$\chi^2(6)$ =	7.60	[0.27]
<u>Univariate tests:</u>			
	<i>Dq_t</i>	<i>Dp_t</i>	<i>Dr_t</i>
Autocorrelation order 2, F(2,29):	0.50 [0.61]	0.06 [0.95]	1.80 [0.18]
Autocorrelation order 4, F(4,27):	1.34 [0.28]	0.71 [0.59]	1.15 [0.35]
Autocorrelation order 6, F(6,25):	1.14 [0.37]	0.56 [0.76]	0.94 [0.49]
ARCH (1), F(1,29):	0.15 [0.71]	1.36 [0.25]	0.02 [0.89]
Heteroskedast. (squares), F(26,4):	0.23 [0.99]	0.07 [1.00]	0.54 [0.85]
Normality, $\chi^2(2)$:	2.58 [0.27]	5.43 [0.07]	9.18 [0.01] *
<u>Goodness-of-fit:</u>			
r	0.98	1.00	0.81
σ_ε	0.184	0.018	0.012

Note: The VAR model has a lag length of 4 (k=4). The F-tests are small sample approximations to Lagrange Multiplier tests, being adjusted for degrees of freedom. Normality test of Doornik and Hansen (1994). For a description of the tests, see Doornik and Hendry (1997). Numbers in brackets are critical significance levels. * and ** indicate misspecification at the 5% and 1% significance level, respectively. r is the correlation between actual and fitted values for each equation (variables in levels). σ_ε is the standard deviation of the residual term.

Table 3. Cointegration Analysis in the VAR Model
Estimation sample 1952-1996

<u>Cointegrating rank:</u>			
Rank(Π) ($r =$)	0	1	2
Eigenvalue	0.54	0.36	0.17
Trace test ¹⁾	63.3 ***	28.3 **	8.2
Trace test (adj. for df.) ²⁾	46.4 **	20.7	6.0
95 % critical test value	42.2	25.5	12.4
97.5 % critical test value	45.0	27.9	14.1
99 % critical test value	48.6	30.7	16.4
<u>Standardized eigenvectors β^*:</u>			
	b_1^*	b_2^*	b_3^*
q_t	1.000	0.066	0.031
p_t	-1.037	1.000	0.008
r_t	5.423	-15.338	1.000
t	-0.011	-0.053	-0.004
<u>Standardized loadings α: ³⁾</u>			
	$b_1^* \alpha_t^*$	$b_2^* \alpha_t^*$	$b_3^* \alpha_t^*$
Δq_t	-0.877	-0.136	8.049
Δp_t	-0.017	-0.051	-0.565
Δr_t	-0.017	0.010	-0.593

Note: Maximum Likelihood Estimation by the Johansen-method, cf. Johansen (1996). The trace tests test for each value of r the null hypothesis H_0 : $\text{rank}(\Pi) \leq r$ against the alternative H_A : $\text{rank}(\Pi) > r$. The null is rejected iff the trace statistic is larger than the critical test value. Critical values from Table 15.4 in Johansen (1996). *, ** and *** indicate rejection of the null at the 5%, 2.5% and 1% significance level, respectively. The standardized eigenvectors are normalized on the diagonal wrt. the endogenous variables. Corresponding I-loadings.

¹⁾ The asymptotic trace test of the Johansen-method.

²⁾ Small sample approximation to the asymptotic trace test, obtained by adjusting for degrees of freedom, cf. Reimers (1992).

³⁾ $X_t^* = (q_t, p_t, r_t, t)'$

Table 4. Estimation and Testing of the Cointegrating Relation: Single-Equation-Analysis

The regression is: (*) $q_t = b_0 + b_1 p_t + b_2 t + b_3 r_t + \sum_{i=-N_1}^{N_1} g_{1i} \Delta p_{t-i} + \sum_{i=-N_2}^{N_2} g_{2i} \Delta r_{t-i} + \sum_{i=1}^{N_3} f_i ecm_{t-i} + n_t$, where $ecm_t \equiv q_t - b_0 - b_1 p_t - b_2 t - b_3 r_t$

(*) is the static cointegrating regression augmented by current, leaded and lagged first differences of p_t and r_t and lagged error correction terms ecm_t . The augmenting terms in each regression are indicated in the first column.

Regression	Sample (no. obs.) No. of regressors	Coefficient Estimates (standard errors)				t-test on price level coeff. (critical sign. level) ¹⁾	
		β_0	β_1	β_2	β_3	$H_0: \beta_1=0$	$H_0: \beta_1=1$
1. No augmentation	1948-1996 (49) 4	0.372 (0.99)	0.898 (0.22)	0.024 (0.014)	-11.32 (1.80)	-	-
2. 1 lead, 1 lag and current first differences of p_t and r_t ($N_1=N_2=1$, $N_3=0$ in (*))	1950-1995 (46) 10	0.283 (0.60)	0.948 (0.13)	0.015 (0.009)	-5.04 (2.15)	7.18 * (0.000)	0.39 (0.697)
3. 2 leads, 2 lags and current first differences of p_t and r_t ($N_1=N_2=2$, $N_3=0$ in (*))	1951-1994 (44) 14	0.001 (1.18)	1.024 (0.26)	0.012 (0.016)	-7.20 (4.35)	3.94 * (0.000)	0.09 (0.928)
4. Δp_{t-1} , Δr_{t+1} and Δr_{t+2}	1950-1994 (45) 7	0.464 (0.83)	0.913 (0.18)	0.024 (0.012)	-11.99 (1.90)	5.02 * (0.000)	0.48 (0.631)
5. Δp_{t-1} , Δr_{t+2} and ecm_{t-1} (NLS)	1950-1994 (45) 7	-0.003 (1.31)	1.007 (0.29)	0.016 (0.019)	-9.04 (1.68)	3.46 * (0.001)	0.02 (0.984)

Note: Entry 1 shows the results from estimating (OLS) the static cointegrating regression (no augmentation), including indicative OLS standard errors. Entries 2 through 4 give the results from the Phillips and Loretan (1991) OLS procedure using different augmentations (as tabulated). In entry 2 the standard errors of the coefficient estimates are adjusted to take account of AR(5) serial correlation in the disturbance term (v_t), using the method suggested by Hamilton (1994, p. 608f). Standard errors in entries 3 and 4 are OLS standard errors as the disturbance term fulfills the white noise requirements. Entry 5 uses the Phillips and Loretan (1991) NLS procedure. NLS standard errors, calculated from numerical derivatives of the sum of squared residuals.

¹⁾ Critical significance level for two-sided t-test, calculated from standard normal distribution (asymptotic test). A ‘*’ indicates that the null is rejected at the 5% significance level.

Table A.1. Cointegration Analysis: Estimates and Tests

Single-equation cointegration analysis following the Engle and Granger (1987) two-step procedure. Residuals-based tests for cointegration.

In the cointegrating regression, stock prices (q_t) are regressed on a constant term (const), measures for the general price level and real production, and the discount rate proxy (r_t), cf. (5). Sample 1948-1996

Model	Candidate cointegrating vector [OLS standard errors] 1)	CRDF (no. of lags) 2)	Critical test value at significance level 3)		Test conclusion: Cointegration at 10% significance level ?
			10%	5%	
$q_t, \text{const}, \text{pyf}_t, \text{yf}_t, r_t$	(1 ; -1.67 ; -1.38 ; 0.16 ; 11.17) [0 ; 0.65 ; 0.14 ; 0.28 ; 2.10]	-2.834 (0)	-3.990 (N=4, No Trend, T=48)	-4.338	No
$q_t, \text{const}, \text{pyfi}_t, \text{yfi}_t, r_t$	(1 ; 0.02 ; -1.41 ; -0.13 ; 8.78) [0 ; 0.29 ; 0.08 ; 0.13 ; 1.76]	-3.717 (0)	-3.990 (N=4, No Trend, T=48)	-4.338	No
$q_t, \text{const}, \text{yf}_t^*, r_t$	(1 ; -3.38 ; -0.87 ; 12.98) [0 ; 0.11 ; 0.04 ; 2.32]	-3.162 (1)	-3.586 (N=3, No Trend, T=47)	-3.927	No
$q_t, \text{const}, p_t, t, r_t$	(1 ; -0.37 ; -0.90 ; -0.02 ; 11.32) [0 ; 0.99 ; 0.21 ; 0.01 ; 1.80]	-4.163 (1)	-4.032 (N=3, With Trend, T=47)	-4.381	Yes
$q_t, \text{const}, \text{pyf}_t, t, r_t$	(1 ; -2.56 ; -0.79 ; -0.03 ; 11.74) [0 ; 0.57 ; 0.23 ; 0.01 ; 1.90]	-3.920 (1)	-4.032 (N=3, With Trend, T=47)	-4.381	No
$q_t, \text{const}, \text{pyfi}_t, t, r_t$	(1 ; -1.29 ; -1.10 ; -0.02 ; 9.18) [0 ; 0.52 ; 0.18 ; 0.01 ; 1.59]	-4.460 (1)	-4.032 (N=3, With Trend, T=47)	-4.381	Yes

Note: For definition of variables entering the model, see text.

1) The candidate cointegrating vector is normalized on stock prices. In terms of (5), the vector is $(1; -\hat{\theta}_0; -\hat{\theta}_1; -\hat{\theta}_2; -\hat{\theta}_3)$. OLS standard errors are indicative only.

2) CRDF is the Dickey and Fuller (1979) t-test statistic of the null of a unit root in the OLS residuals from the cointegrating regression. The null is rejected in favor of the stationary alternative (and cointegration is accepted) if the test statistic is negative and larger in absolute value than the critical test value. Number of augmenting lags of the first differences of the OLS residuals used in the unit root regression shown in parenthesis.

3) Small sample critical values from MacKinnon (1991). N=number of I(1) variables in the model; 'No Trend' and 'With Trend' indicate whether a deterministic trend is included in the cointegrating regression; T=number of observations in the unit root regression, cf. Table 1 in MacKinnon (1991).

Figure B1. Stock Return and Inflation

Annual stock return and inflation for the (by row) 1-, 5- and 10-year horizon.
Sample periods 1949-1996, 1953-1996 and 1958-1996, respectively.




Table B1.a. Phillips and Perron (1988) Z_t -Test for Unit Root

Series:	<i>Lag length (l)</i>						
	0	1	2	3	4	5	6
<u>1-Year Horizon, 1949-1996:</u>							
$S1_t$	-8.52***	-8.52***	-8.53***	-8.55***	-8.61***	-8.67***	-8.72***
$I1_t$	-2.79*	-2.94**	-2.80*	-2.70*	-2.73*	-2.78*	-2.86*
$GR1_t$	-8.14***	-8.15***	-8.25***	-8.41***	-8.69***	-9.00***	-9.45***
$GY1_t$	-5.27***	-5.24***	-5.24***	-5.32***	-5.41***	-5.50***	-5.55***
<u>5-Year Horizon, 1953-1996:</u>							
$S5_t$	-3.16**	-3.11**	-3.23**	-3.22**	-3.28**	-3.15**	-3.08**
$I5_t$	-0.31	-0.63	-0.78	-0.91	-1.01	-1.08	-1.13
$GR5_t$	-3.25**	-3.35**	-3.41**	-3.36**	-3.31**	-3.11**	-2.96**
$GY5_t$	-1.28	-1.41	-1.43	-1.49	-1.56	-1.54	-1.49
<u>10-Year Horizon, 1958-1996:</u>							
$S10_t$	-2.19	-2.15	-2.18	-2.19	-2.18	-2.20	-2.20
$I10_t$	-0.25	-0.57	-0.74	-0.88	-0.99	-1.08	-1.16
$GR10_t$	-2.80*	-2.81*	-2.80*	-2.78*	-2.79*	-2.82*	-2.82*
$GY10_t$	-0.28	-0.40	-0.41	-0.49	-0.58	-0.64	-0.68
<u>Critical test values:</u>			10 %	5 %	2.5 %	1 %	
Without trend			-2.60	-2.93	-3.22	-3.58	

Note: See note to Table 1.a. All regressions include a constant term, while no trend is allowed for. *, ** and *** denote rejection of the null of a unit root at the 10%, 5% and 1% significance level, respectively.

Table B1.b. Kwiatkowski *et al.* (1992) Test for Unit Root

Series:	<i>Lag length (l)</i>						
	0	1	2	3	4	5	6
<u>1-Year Horizon, 1949-1996:</u>							
S1 _t	0.12	0.16	0.16	0.17	0.19	0.20	0.21
I1 _t	0.61**	0.36*	0.28	0.23	0.20	0.17	0.16
GR1 _t	0.05	0.06	0.07	0.09	0.11	0.12	0.15
GY1 _t	1.05***	0.84***	0.74***	0.64**	0.56**	0.51**	0.48**
<u>5-Year Horizon, 1953-1996:</u>							
S5 _t	0.42*	0.26	0.20	0.18	0.16	0.16	0.16
I5 _t	1.01***	0.52**	0.36*	0.28	0.23	0.20	0.18
GR5 _t	0.32	0.20	0.17	0.15	0.16	0.17	0.19
GY5 _t	2.43***	1.28***	0.90***	0.70**	0.59**	0.52**	0.47**
<u>10-Year Horizon, 1958-1996:</u>							
S10 _t	1.31***	0.74***	0.54**	0.44*	0.38*	0.34*	0.31*
I10 _t	1.19***	0.61**	0.42*	0.33	0.27	0.24	0.21
GR10 _t	0.73**	0.44*	0.34	0.29	0.26	0.24	0.22
GY10 _t	3.04***	1.57***	1.07***	0.83***	0.68**	0.58**	0.52**
<u>Critical test values:</u>							
			10 %	5 %	1 %		
Without trend			0.35	0.46	0.74		

Note: See note to Table 1.b. No trend is allowed for in the tests, i.e., the null hypothesis is mean-stationarity. *, ** and *** denote rejection of the null (i.e., a unit root is present) at the 10%, 5% and 1% significance level, respectively.

Table B2. Return Regressions for the 1-, 5- and 10-Year Investment Horizon.

The model is: (*) $Sk_t = b_0 + b_1Ik_t + b_2GYk_t + b_3GRk_t + e_t$, $k = 1, 5, 10$ years

The 'simple' model in the table only includes inflation as an explanatory variable ($\beta_2 = \beta_3 = 0$), while the 'extended' model includes all variables in (*). In the 'reduced extended' model, we have removed the insignificant variables from the extended model.

Horizon (years)	Sample (sample size)	Model	Coefficient estimates (standard errors)				Goodness-of-fit-test on inflation coeff. 2)		(critical sign. level) 3)	
			β_0	β_1	β_2	β_3	R^2	σ_e	$H_0: \beta_1=0$	$H_0: \beta_1=1$
1	1949-1996 (48)	Simple 1)	0.102 (0.029)	-0.106 (0.61)	-	-	0.57	0.18	-0.17 (0.865)	-
		Extended 1)	0.152 (0.045)	-0.205 (0.68)	-1.02 (0.88)	-5.63 (1.8)	0.67	0.16	-0.30 (0.764)	-
		Reduced extended 1)	0.111 (0.032)	-0.0368 (0.73)	-	-6.01 (1.7)	0.66	0.16	-0.05 (0.959)	-
5	1953-1996 (44)	Simple	0.0617 (0.025)	0.936 (0.61)	-	-	0.11	0.073	1.53 (0.126)	-
		Extended	0.0908 (0.041)	0.936 (0.57)	-0.702 (0.69)	-7.27 (2.8)	0.34	0.064	1.64 (0.101)	-
		Reduced extended	0.0643 (0.025)	1.01 (0.55)	-	-7.35 (2.8)	0.33	0.064	1.84 (0.066)	0.03 (0.979)
10	1958-1996 (39)	Simple	0.0665 (0.014)	0.858 (0.28)	-	-	0.21	0.041	3.06 * (0.002)	-0.51 (0.610)
		Extended	0.109 (0.031)	0.808 (0.28)	-0.944 (0.55)	-8.85 (2.4)	0.59	0.030	2.89 * (0.004)	-0.68 (0.497)
		Reduced extended	0.0668 (0.014)	1.01 (0.21)	-	-9.99 (2.6)	0.54	0.031	4.81 * (0.000)	0.04 (0.967)

Note: All coefficients are estimated by OLS. Standard errors of the coefficient estimates are OLS errors for the 1-year horizon, respectively Newey and West (1987) errors for the 5- and 10- year horizons. The reported Newey and West (1987) errors allow for heteroskedasticity and serial correlation in the disturbance term (e_t) up to lag 5.

1) Impulse dummies included for 1972 and 1983.

2) Not corrected for serial correlation or heteroskedasticity in the disturbance term.

3) Critical significance level for two-sided t-test, calculated from standard normal distribution (asymptotic test). Based on Newey and West (1987) standard errors for the 5- and 10-year horizons. A '*' indicates that the null is rejected at the 5% significance level.

Figure B2. Recursive Estimation of the 10-Year Return Regression

Recursive point estimates (solid line) and 95% confidence bands for the coefficients of the 10-year return regression (reduced extended model). Recursive least squares. Full sample: 1958-1996.

