Inflation prediction from the term structure: the Fisher equation in a multivariate SDF framework

Chiona Balfoussia

University of York

Mike Wickens

University of York and CEPR

December 2004

Preliminary version

Abstract

In this paper we propose a new way to extract inflation information from the term structure. We rehabilitate the workhorse of inflation prediction, the Fisher equation, by setting it in the context of the stochastic discount factor (SDF) asset pricing theory. We propose a multivariate estimation framework which models the term structure of interest rates in a manner consistent with the SDF theory while at the same time generating and including an often omitted time varying risk component in the Fisher equation. The joint distribution of excess holding period bond returns of different maturity and fundamental macroeconomic factors is modelled on the basis of the consumption CAPM, using multivariate GARCH with conditional covariances in the mean to capture the term premia. We apply this methodology to the U.S. economy and re-examine the Fisher equation at horizons of up to one year. We find it offers substantial evidence in support of the Fisher equation and greatly improves its goodness of fit, at all horizons.

Keywords: Inflation, Fisher equation, Term structure, the stochastic discount factor model, term premia, GARCH

JEL Classification: G1, E4, C5

Acknowledgements: We would like to express particular thanks to Peter Smith for many helpful suggestions.

1 Inflation Prediction and the Term Structure of Interest Rates

Macroeconomists and financial analysts rarely find that they have a lot to discuss. All too often it seems as if their perspectives on the same economic phenomena are vastly different. However, the term structure of interest rates is one of the few subjects which link the academic fields of macroeconomics and finance. Depending on one's point of view, the level and slope of the yield curve can be indicators of the current stance of monetary policy, predictors of future movements in real output and inflation or reflections of the market's assessment of the risk and expected return of bonds of different maturities. It is precisely this plethora of interpretations and the potential wealth of hidden information which makes the term structure of interest rates fascinating and never out of fashion. As a result, the main body of the term structure literature is, in one way or another, concerned with the informative content of bonds and ways to extract and exploit it.

Yields, spreads and forward rates have all been examined rigorously to uncover any predictive power that may be hidden in them. Inflation prediction forms a substantial part of this literature. The theoretical modelling and successful prediction of inflation is one of the most consequential issues in macroeconomics. In recent years, central banks have gradually abandoned the control of monetary aggregates and have directly or indirectly implemented inflation targeting rules, often by adopting aggressive Taylor rules. Indeed, the rate of inflation is today considered the conventional final target of monetary policy in most developed countries. Forecasting inflation has hence become increasingly important for the conduct of monetary policy and the assessment of central bank credibility. Furthermore, it has also become crucial to private agents who try to understand and react to the central bank's behaviour. Hence, more and more time and effort is being channeled towards capturing and forecasting inflation dynamics.

Various models and techniques have been proposed to recover information on the future path of inflation. However, the literature seems to be inconclusive as to which is the optimal predictive method. The information content in the dynamics of past inflation seems to suffice to predict future inflation and few other variables are able to offer some improvement on this prediction (one example is Cecchetti, Chu and Steindel(00)). In fact, inflation seems to be very well described even by a simple AR(1) process, often a near random walk. Canova (01) compares the forecasting performance of some leading models of inflation for the cross-section of G-7 countries and shows that these simple, univariate models can often outperform the bivariate and trivariate models suggested by economic theory. However, despite their good fit, there is no theoretical justification for such *ad hoc* statistical specifications. Hence we shall instead focus entirely on theoretical models of inflation prediction, and in particular on the theory developed to model and test the predictive power of the term structure of interest rates.

Under appropriate assumptions about the behaviour of financial markets, the term structure of interest rates is thought to reflect market expectations of risk premia, real interest rates and future inflation. This academic conviction is also reflected in policy making, where the forward-looking nature of the yield curve has meant it is one of the financial market indicators typically monitored by monetary authorities and discussed in the context of monetary policy debate. If markets are rational and efficient, their expectations should be unbiased predictors. If they have more information than policy makers or process information more efficiently, then it would be instructive for the monetary authority to extract the market expectations of future inflation implicit in the term structure and use them in the formulation and adjustment of monetary policy. They can serve as a useful check on the reliability of forecasts obtained from a full general equilibrium model of an economy. Even if markets are less than fully efficient, market participants and economic agents will still base their economic decisions on their own, *albeit* imperfect, forecasts. Market expectations of inflation would therefore reveal investors' confidence in and reaction to policies and news, whatever the degree of efficiency of the market. They are hence a valuable piece of information.

The extent and sophistication of the use of the yield curve for inflation prediction varies from one country to another. In the U.S., the term structure has, at times, been an explicit factor taken into account in setting policy. However, the yield curve is most commonly included in monetary policy discussions as a simple leading indicator of inflation, and often in a rather informal manner, reflecting the relevant academic evidence.

To date, a relatively extensive part of the academic literature has theorised and formally examined the informational and predictive content of the term structure with regard to the rate of inflation. Work on this area started in the early 1980's (see for example Fama and Gibbons 84) while recent contributions have been made, among others, by Schich (99), Stock and Watson (99) and Estrella, Rodrigues and Schich (00). The majority of the work in this area makes two assumptions about the economy and the financial world; that the underlying real interest rate is constant through time and that the risk premium included in the term structure of interest rates is also constant, if not zero. Balfoussia and Wickens (04) have shown that not only is the excess holding-period return risk premium (and hence the rolling risk premium) non-zero, but neither is it constant, nor indeed is the slope of its term structure constant through time. To the contrary, it seems to vary dramatically under different monetary policies and levels of perceived macroeconomic risk. The assumption of its constancy is hence unfounded. This is confirmed by numerous papers, where the inclusion of a rolling risk premium proxy in univariate estimation of the Fisher equation proves significant. This paper builds on the Balfoussia and Wickens (04) paper and on the existing literature on inflation prediction by re-examining its principal theoretical tool, the Fisher equation, in view of our findings on the term structure risk premium, in a multivariate SDF methodology.

The paper is structured as follows. In Section 2 we construct the theoretical model proposed.

We first discuss the Fisher equation and the SDF asset pricing framework in some detail. Subsequently, we propose a multivariate estimation framework, which models the term structure of interest rates in a manner consistent with the SDF theory, while at the same time generating and including a time varying risk component in the Fisher equation. The econometric methodology, set out in Section 3, modifies the multi-variate GARCH model to fit our specification. The data are described in Section 4. Section 5 presents the estimates obtained when this methodology is applied to the U.S. economy, re-examining the Fisher equation at horizons of up to one year. Finally, in Section 6 we present our conclusions.

2 Theoretical Framework

2.1 Basic concepts

We use the following notation. $P_{n,t}$= price of an n-period, zero-coupon (pure discount), default-free, bond at t, where $P_{0,t}$= 1 as the pay-off at maturity is 1. $R_{n,t}$= yield to maturity of this bond, where the one-period, risk-free rate $R_{1,t}$= s_t. $h_{n,t+1}$ = return to holding an n-period bond for one period from t to $t+1$. It follows that

$$P_{n,t} = \frac{1}{[1 + R_{n,t}]^n}$$

and

$$1 + h_{n,t+1} = \frac{P_{n-1,t+1}}{P_{n,t}} \tag{1}$$

If $p_{n,t} = \ln P_{n,t}$ then, taking logs

$$h_{n,t+1} \simeq p_{n-1,t+1} - p_{n,t} = nR_{n,t} - (n-1)R_{n-1,t+1} \tag{2}$$

In the absence of arbitrage opportunities, after adjusting for risk, investors are indifferent between holding an n-period bond for one period and holding a risk-free 1-period bond. In the absence of default, the risk is due to the price of the bond next period being unknown this period.

$$E_t[h_{n,t+1}] = s_t + \rho_{n,t} \tag{3}$$

where $\rho_{n,t}$ is the risk, or term, premium on an n-period bond at time t.

2.2 A general equilibrium SDF model of the term structure

The SDF model relates the price of an n-period zero-coupon bond in period t to its discounted price in period $t+1$ when it has $n-1$ periods to maturity. Thus

$$P_{n,t} = E_t[M_{t+1}P_{n-1,t+1}]$$

where M_{t+1} is a stochastic discount factor, or pricing kernel. It follows that

$$E_t[M_{t+1}(1 + h_{n,t+1})] = 1$$

and for $n = 1$,

$$(1 + s_t)E_t[M_{t+1}] = 1$$

If $P_{n,t}$ and M_{t+1} are jointly log-normally distributed and $m_{t+1} = \ln M_{t+1}$ then

$$p_{n,t} = E_t(m_{t+1}) + E_t(p_{n-1,t+1}) + \frac{1}{2}V_t(m_{t+1}) + \frac{1}{2}V_t(p_{n-1,t+1}) + Cov_t(m_{t+1}, p_{n-1,t+1}) \tag{4}$$

and as $p_{o,t} = 0$,

$$p_{1,t} = E_t(m_{t+1}) + \frac{1}{2}V_t(m_{t+1}) \tag{5}$$

Subtracting 5 from 4 and re-arranging gives the no-arbitrage equation

$$E_t(p_{n-1,t+1}) - p_{n,t} + p_{1,t} + \frac{1}{2}V_t(p_{n-1,t+1}) = -Cov_t(m_{t+1}, p_{n-1,t+1}) \tag{6}$$

Using 1 this can be re-written in terms of holding period returns as

$$E_t(h_{n,t+1} - s_t) + \frac{1}{2}V_t(h_{n,t+1}) = -Cov_t(m_{t+1}, h_{n,t+1}) \tag{7}$$

This is the fundamental no-arbitrage condition for an n-period bond[1] , which must be satisfied to ensure there are no arbitrage opportunities across the term structure. The term on the right-hand side is the term premium and $\frac{1}{2}V_t(h_{n,t+1})$ is the Jensen effect. Comparing equations 3 and 7, we note the SDF model implies that

$$\rho_{n,t} = -\frac{1}{2}V_t(h_{n,t+1}) - Cov_t(m_{t+1}, h_{n,t+1})$$

Empirical work on the term structure can be distinguished by the choice of $\rho_{n,t}$ and the discount factor m_t. The expectations hypothesis assumes that $\rho_{n,t} = 0$ but a vast amount of evidence in the literature rejects this assumption. Hence, we shall assume that the risk premium is non-negligible and hence must be explicitly modelled.

To complete the specification, it is necessary to also specify m_t. Assuming joint log-normality of the excess returns and the factors, m_t is a linear function of some underlying factors. However, the SDF model does not specify which factors should be used. For the purpose of this paper, we shall use observable factors, drawing on a general equilibrium model for guidance as to which these should be.

2.3 The Fisher Equation

We now take a closer look at the Fisher equation, the main theoretical proposition under examination in this work, with an aim to incorporate it in a multivariate estimation framework, within the SDF methodology. The Fisher equation, first made by Fisher in 1930, is a simple and intuitive relationship, which decomposes the nominal interest rate of a given maturity into a real rate and an inflation expectations component, both for the period from the present to the maturity of the instrument. For the short rate s_t it takes the following, familiar form.

$$s_t = r_t + E_t[\pi_{t+1}] \quad (8)$$

[1] Arbitrage opportunities are excluded if and only if there exists a *positive* stochastic discount factor M_{t+1} that prices *all* assets (Cochrane 2001). In the models described in this section, a positive discount factor is used to price bonds of all maturities.

where r_t is the one-period underlying real interest rate and $E_t[\pi_{t+1}]$ is the one period ahead inflation expectation, conditional on information available at time t.[2]

From equations 2 and 3 it follows that

$$E_t[h_{n,t+1}] = E_t[nR_{n,t} - (n-1)R_{n-1,t+1}] = s_t + \rho_{n,t} \tag{9}$$

Hence

$$R_{n,t} = \frac{n-1}{n} E_t[R_{n-1,t+1}] + \frac{1}{n}(s_t + \rho_{n,t})$$

and, by forward substitution, the yield of a nominal bond can be expressed as

$$\begin{aligned} R_{n,t} &= \frac{1}{n}\sum_{i=0}^{n-1} E_t(s_{t+i} + \rho_{n-i,t+i}) \\ &= \frac{1}{n}\sum_{i=0}^{n-1} E_t s_{t+i} + \frac{1}{n}\sum_{i=0}^{n-1} E_t \rho_{n-i,t+i} \end{aligned} \tag{10}$$

By substituting forward equation 8 into 10 we obtain the following expression

$$\begin{aligned} R_{n,t} &= \frac{1}{n}\sum_{i=0}^{n-1} E_t r_{t+i} + \frac{1}{n}\sum_{i=1}^{n} E_t \pi_{t+i} + \frac{1}{n}\sum_{i=0}^{n-1} E_t \rho_{n-i,t+i} \\ &= \frac{1}{n}\sum_{i=0}^{n-1} E_t r_{t+i} + \frac{1}{n}\sum_{i=1}^{n} E_t \pi_{t+i} + \omega_{n,t} \end{aligned} \tag{11}$$

where $\omega_{n,t} = \frac{1}{n}\sum_{i=0}^{n-1} E_t \rho_{n-i,t+i}$ is the rolling risk premium of an n-period bond at time t. Hence we can decompose each yield into three components: a real, a nominal and a risk component, all of which are functions of the time to maturity of the bond. More specifically, each of these components is the average of the expected values of the corresponding variable from time t to the maturity of the bond. They are hence all forward-looking. For example, the inflation component is the average expected inflation over the life of the bond.

Inflation prediction from the term structure of interest rates is essentially an effort to recover the forward-looking inflation component in equation 11. The success of this effort depends on our

[2] It is worth mentioning that this form of the Fisher equation is accurate only for the short rate s_t which, by assumption, is taken to be risk free. The same cannot be assumed to hold for longer maturities. This is the reason we substitute the 1-period Fisher equation into 10, rather than simply taking the n-period form of the Fisher equation.

ability to disentangle the inflation component from the real and risk components, and from noise in the yield. However, extracting expectations is, by any measure, not a trivial problem.

Solving equation 11 for the expected inflation component we came to the following form of the equation

$$E_t \pi_{n,t} = R_{n,t} - \frac{1}{n} \sum_{i=0}^{n-1} E_t r_{t+i} - \omega_{n,t} \tag{12}$$

where $E_t \pi_{n,t} = \dfrac{1}{n} \sum_{i=1}^{n} E_t \pi_{t+i}$. Hence $\pi_{n,t}$ is the average, realised inflation between times t and $t+n$, where n denotes the horizon over which inflation is predicted as well as the maturity of the corresponding bond. If market expectations are rational, expected average inflation will differ from actual inflation only by an unpredictable noise term. This is the form which will be examined and tested in this paper, as opposed to the slightly different form which is usually used in the literature. They are both essentially identical. However the usual form describes the expected change in inflation between two points in the future, while equation 12 describes the expected change in the price level between the current period and a future period. In other words, it describes the forward looking inflation expectations formed today for a given horizon as of today.

Typically, the risk premium is either omitted from the theoretical analysis and hence effectively set to zero, implying risk neutrality, or it is explicitly assumed to be time-invariant and hence thought to be captured by the constant term in the regression. As a result, the information content of the term structure on future inflation is often defined simply as the ability of the slope of the term structure to predict changes in inflation rates. Indeed, if both real interest rates and risk premia were constant over time, expected inflation over a given horizon would simply be a linear function of the nominal average interest rate over the same period. However, having already established that the risk premium is not only a significant determinant of the term structure but also significant in the Fisher equation, we cannot omit it in either specification, as that would lead to biased estimates. To address this issue, we propose a multivariate framework within which to account for the risk premium component, while jointly estimating both the Fisher equation and the term structure equations.

2.4 The Fisher Equation in a SDF Model

Our aim in this paper is to examine whether the omission of a risk premium term could account for the low predictive power of the term structure, when using models based on the Fisher equation. Hence we will test whether the inclusion of a risk premium term in this standard inflation prediction model leads to an improvement of its predictive power. However, we shall do so by jointly modeling inflation and the term structure, in a multivariate, general equilibrium SDF framework. This allows us to estimate the risk premium directly from the term structure and to contemporaneously include it in the Fisher equation, without needing to resort to a two-step method and univariate estimation.

Different formulations of the SDF model for the term structure are discussed and tested in Balfoussia and Wickens (04). In the C-CAPM general equilibrium framework, the log-stochastic discount factor is

$$m_{t+1} \simeq \ln \beta - \sigma_t \frac{\Delta C_{t+1}}{C_t} - \pi_{t+1} \tag{13}$$

where π_{t+1} is the rate of inflation. Thus the C-CAPM implies that the stochastic discount factor depends on the growth rate of consumption and the rate of inflation. Substituting this result into the corresponding no-arbitrage condition they obtained

$$E_t(h_{n,t+1} - s_t) = -\frac{1}{2}V_t(h_{t+1}) + \sigma_t Cov_t(\frac{\Delta C_{t+1}}{C_t}, h_{n,t+1}) + Cov_t(\pi_{t+1}, h_{n,t+1}) \tag{14}$$

where the right hand side of this expression is the risk premium $\rho_{n,t}$ of any bond at time t (assuming log-normality of returns) adjusted by a Jensen term. This implies that the greater the predicted covariation of the risky return with consumption growth and inflation, the higher the risk premium. In other words, assets are being priced in accordance to the insurance they offer against adverse movements in consumption. This expression provides a no-arbitrage condition for each bond individually, as well as for the bond market as a whole, as it eliminates arbitrage possibilities across the entire term structure of interest rates.

Equation 14 can be used to test the validity of the C-CAPM. However, our specification in this work shall not be exactly that of equation 14. Our focus is on modeling and explaining inflation, rather that on the term structure itself. Hence, we need to specify the stochastic discount factor slightly differently if we are to use the Fisher equation as our inflation specification. Nevertheless, we will be guided by the economic intuition of the CAPM in formulating our SDF. The stochastic discount factor shall, as implied, be a function of inflation and consumption. However, since we want to fit future inflation for horizons of up to one year, we shall this time use realised, *ex post* inflation, over different horizons in the future in the SDF, rather than lagged, one month inflation. The log stochastic discount factor will be specified as follows. For the n-month horizon

$$m_{n,t+1} \simeq a_n + \beta_n \frac{\Delta C_{t+1}}{C_t} + \gamma_n \pi_{t+n} \qquad (15)$$

where π_{t+n} denotes realised inflation between period t and period $t+n$.

Although still a linear function of inflation and consumption, this is not exactly the stochastic discount factor specification implied by C-CAPM. However, it remains in line with the intuition of C-CAPM, while offering the advantage of allowing us to jointly model inflation over an n-period horizon, in our setup. This also means the holding-period return risk premia generated include information on inflation over an n-period horizon. This risk premium is endogenously generated and it will form the basis of the risk component to be included in the Fisher equation.

Since we have not exactly specified the stochastic discount factor as in equation 13 and, furthermore, since in Balfoussia and Wickens (04) the C-CAPM, was rejected against alternative hypotheses, we shall neither constrain the coefficients of the conditional covariances in the term structure equations to the theoretical values implied by C-CAPM, nor constrain them to be equal across maturities. In other words, we shall essentially be using a general excess holding-period return specification. If $z_{i,t+1}$ are the factors which compose the pricing kernel, this specification is

$$E_t(h_{n,t+1} - s_t) = -\frac{1}{2}V_t(h_{n,t+1}) - \sum_i b_{n,i}Cov_t(h_{n,t+1}, z_{i,t+1}) \tag{16}$$

This formulation places no restrictions on the coefficients of the conditional covariance terms, and is the one which will be used in this paper, to model the term structure.

Since s_t is known at time t, we can, as before, express equation 16 in terms of excess holding-period returns alone

$$E_t(h_{n,t+1} - s_t) = -\frac{1}{2}V_t(h_{t+1} - s_t) - \sum_i b_{n,i}Cov_t(h_{t+1} - s_t, z_{i,t+1}) \tag{17}$$

Hence, assuming the stochastic discount factor is specified as in 15, we obtain the exact excess return specification for the n-period horizon

$$E_t(h_{n,t+1} - s_t) = -\frac{1}{2}V_t(h_{t+1} - s_t) + b_{n,1}Cov_t(\frac{\Delta C_{t+1}}{C_t}, h_{n,t+1} - s_t) + b_{n,2}Cov_t(\pi_{n,t+1}, h_{n,t+1} - s_t) \tag{18}$$

where the right-hand side of the expression is a measure of the risk premium, $\rho_{n-i,t+i}$ adjusted by a Jensen term.

We now return to our main variable of interest, inflation, which will be specified as the Fisher equation. The Fisher equation for an n-month horizon is

$$E_t\pi_{n,t} = R_{n,t} - \frac{1}{n}\sum_{i=0}^{n-1} E_t r_{t+i} - \omega_{n,t}$$

where $\omega_{n,t} = \frac{1}{n}\sum_{i=0}^{n-1} E_t\rho_{n-i,t+i}$. Hence, for the three-month horizon one would want to be able to include estimates of $\rho_{3,t+i}$, $\rho_{2,t+i}$ and $\rho_{1,t+i}$ in the Fisher equation, having obtained them using the SDF specification 13, i.e. month-on-month inflation. However, generating each of these components in a multivariate framework, for any horizon, would require the estimation of very large systems, and the availability of data at every period, i.e., in our case monthly data for bonds of every maturity (i.e. every month) in the horizon examined. Both make the exact estimation of this equation impossible. Having established from our work so far that the $\rho_{n-i,t+i}$ for different n are highly correlated and having modified our specification to incorporate inflation expectations over the horizon of interest, we can instead try to proxy the rolling risk premium $\omega_{n,t}$ included in the Fisher equation with a function $\tilde{\rho}_{n,t}$ of the excess holding-period return risk premium of the n-period bond alone, obtained using the SDF specification of equation 15 and hence realised *ex post* inflation over n periods ahead $\pi_{n,t+1}$ rather than month-on-month inflation π_{t+1}. The simplest functional form would be to include in the Fisher equation the two covariances of the excess holding-period return of maturity n with the macroeconomic variables, with an aim to proxy the omitted, time-varying rolling risk premium. Hence

$$\tilde{\rho}_{n,t} = \vartheta_{n,1} Cov_t(\frac{\Delta C_{t+1}}{C_t}, h_{n,t+1} - s_t) + \theta_{n,2} Cov_t(\pi_{n,t+1}, h_{n,t+1} - s_t) \tag{19}$$

Using this framework, we can examine the predictive power of the Fisher equation and the contribution of the risk premium proxy terms in particular, by slightly modifying the stochastic discount factor for each estimation in such a way that the length of the horizon is accounted for.

Once again, while explicitly accounting for the risk premium term, we are unable to do the same for the real rate. However, this time we shall not include a constant term to capture the real rate, the reason being that, since the risk premium term is now endogenously generated rather than an exogenously included variable, the inclusion of a constant might also capture a possible constant element of the risk premium itself. Since we cannot identify both the rolling risk premium terms as a function of the conditional covariance of the SDF with the excess holding-period returns and the constant (taken to capture an assumed constant underlying real rate), we chose to leave the real rate to be captured by the error term. Hence, the specification of the Fisher equation to be tested is

$$E_t \pi_{n,t} = R_{n,t} - \tilde{\rho}_{n,t}$$

3 Econometric methodology

Our aim is to estimate the Fisher equation with and without the risk premium term, while jointly estimating the no-arbitrage condition of the term structure. Hence, we must model the joint distribution of the macroeconomic variables, i.e. inflation and consumption, jointly with the excess holding-period returns in such a way that the mean of the conditional distribution of inflation is allowed to include the appropriate rolling risk premium proxy terms. The conditional mean of both inflation and the excess holding-period returns involves selected time-varying second moments of the joint distribution. We therefore require a specification of the joint distribution that admits a time-varying variance-covariance matrix. A convenient choice is the multivariate GARCH-in-mean (MGM) model. For a review of multivariate GARCH models see Bollerslev, Chou and Kroner (1997), and see Flavin and Wickens (1998) for a discussion of specification issues for their use in financial models. We use the vector-diagonal multivariate GARCH-in-mean model, while appropriately adjusting the in-mean equations to our inflation specification.

Let $\mathbf{x}_{t+1} = (h_{n_1,t+1} - s_t,\ h_{n_2,t+1} - s_t, ..., z_{1,t+1},\ z_{2,t+1})'$, where $z_{1,t+1},\ z_{2,t+1}$ are the two macroeconomic variables that give rise to the factors in the SDF model through their conditional covariances with the excess holding-period returns, that is inflation over the selected horizon and real consumption growth. The model can be written

$$\mathbf{x}_{t+1} = \boldsymbol{\alpha} + \Gamma\mathbf{x}_t + \mathbf{B}\mathbf{g}_t + \boldsymbol{\varepsilon}_{t+1}$$

where

$$\boldsymbol{\varepsilon}_{t+1} \mid I_t \sim D[0, \mathbf{H}_{t+1}]$$

$$\mathbf{g}_t = vech\{\mathbf{H}_{t+1}\}$$

The *vech* operator converts the lower triangle of a symmetric matrix into a vector. The distribution is the multivariate log-normal distribution.

Whilst the MGM model is convenient, it is not ideal. It is heavily parameterised which can create numerical problems in finding the maximum of the likelihood function due to the likelihood surface being relatively flat, and hence uninformative. The exact choice of specification of the conditional covariance matrix is hence a compromise between generality and feasibility. The latter requires that we limit the number of parameters to estimate. Accordingly, our specification of $\mathbf{H}_t$ is that of Ding and Engle (1994), the vector-diagonal multivariate GARCH-in-mean

$$\mathbf{H}_t = \mathbf{H}_0(\mathbf{ii}' - \mathbf{aa}' - \mathbf{bb}') + \mathbf{aa}' * \Sigma_{t-1} + \mathbf{bb}' * \mathbf{H}_{t-1}$$

where $\mathbf{i}$ is a vector of ones, $*$ denotes element by element multiplication (the Hadamard product) and $\Sigma_{t-1} = \boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}'_{t-1}$. This is a special case of the diagonal Vech model, in which each conditional covariance depends only on its own past values and on surprises. The restrictions implicitly imposed by this parameterisation of the multivariate GARCH process guarantee positive-definiteness and also substantially reduce the number of parameters to be estimated, thus facilitating computation and convergence. Stationarity conditions are imposed. The estimation is performed using quasi-maximum likelihood.

Inflation is specified following the Fisher equation, as a function of the nominal yield of the corresponding maturity. Hence, for $n = 3, 6, 12$ the inflation specification will be

$$\pi_{n,t} = \delta_{n,1} R_{n,t} + \delta_{n,2}\tilde{\rho}_{n,t} + \varepsilon_{n,t}$$

where, as set out in equation 19, $\tilde{\rho}_{n,t}$ is a function of the conditional covariances generated by the excess holding-period return equation of the maturity corresponding to the chosen inflation prediction horizon n, expected to proxy for the omitted rolling risk premium.

Hence

$$\pi_{n,t} = \delta_{n,1} R_{n,t} + \zeta_{n,1} Cov_t(\frac{\Delta C_{t+1}}{C_t}, h_{n,t+1} - s_t) + \zeta_{n,2} Cov_t(\pi_{n,t+1}, h_{n,t+1} - s_t) + \eta_{n,t} \quad (20)$$

where $\zeta_{n,i} = \delta_{n,2}\vartheta_{n,i}$ for $i = 1, 2$.

We are interested in testing the Fisher equation's validity and predictive power. More specifically, for each horizon n we will test $\delta_{n,1} = 1$ and $\delta_{n,2} \neq 0$ individually and jointly. Since $\delta_{n,2}$ and $\vartheta_{n,i}$ parameters cannot not be separately identified, we shall instead be testing the hypotheses $\delta_{n,1} = 1$ and $\zeta_{n,i} = 0$ for $i = 1$, 2. According to the Fisher equation we expect to accept these hypotheses, and we also expect the inclusion of the proxy $\tilde{\rho}_{n,t}$ to improve the explanatory power of the Fisher equation.

The first 2 equations of the model, the term structure equations, are restricted to satisfy the condition 18. The consumption equation is specified as an AR(1).

4 Data

The complete sample is monthly, from January 1970 to December 1998. Until 1991, the term structure data are those of McCulloch and Kwon.[3] [4] We then use the data extended by Bliss which is based on the same technique. Excess holding-period returns are taken in excess of the one-month risk-free rate provided by K. French[5] [6] . Essentially, this is the one-month Treasury bill rate. Having established that it is necessary to adequately represent the yield curve in order to satisfactorily model the term structure, we include in all estimations not only the bond of maturity equal to the horizon n examined, but also a second one, of medium to long maturity. For the 3-month horizon, bonds of 3 and 24 months to maturity are included; for the 6-month horizon bonds of 6 and 24 months to maturity are included in the estimation while for the 12-

[3] McCulloch used tax-adjusted cubic splines to interpolate the entire term structure from data on most of the outstanding Treasury bills, bonds and notes, see McCulloch (1975) and McCulloch and Kwon (1993). These data have been used extensively in empirical work and are considered very "clean", in that they are based on a broad spectrum of government bond prices and are corrected for coupon and special tax effects. Furthermore, spline-based techniques allow for a high degree of flexibility, since individual curve segments can move almost independently of each other (subject to continuity constraints). Hence they are able to accommodate large variations in the shape of the yield curve.

[4] We are not using the complete term structure dataseries available by McCulloch and Kwon because no real personal consumption data was available for earlier dates.

[5] In constructing holding-period returns we use the change in n-period yields $\Delta R_{n,t+1}$ instead of $R_{n-1,t+1} - R_{n,t}$, since $R_{n-1,t+1}$ is not available. This approximation is very commonplace in the literature, and is likely to be good for medium and long yields. For a discussion on the effects of this approximation see Bekaert, Hodrick and Marshall (1997a).

[6] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/

month horizon the maturities included are of 12 and 60 months. Hence all estimations have four equations, two excess holding-period return equations and two equations for the macroeconomic variables.

The consumption measure used in this work is the month on month growth rate of total personal consumption (TPC). Inflation is the 3, 6 and 12-month ahead realised *ex-post* growth rate of the consumer price index for all urban consumers (CPI3 CPI6 and CPI12 respectively). Our sample has 346 observations for the 3-month horizon, 343 for the 6-month horizon and 337 for the 12-month horizon. All data are expressed in annualised percentages.

Table 3.1 of Chapter 3 presents some descriptive statistics for TPC and for the term structure dataset. The average yield curve is upward sloping. Average excess holding-period returns are positive for all maturities and increase with time to maturity. Like most financial data, excess holding-period returns exhibit excess skewness and kurtosis, particularly for short maturities. Fitting a univariate GARCH(1,1) to them indicates there is also significant heteroskedasticity.[7] Their unconditional variances increase with maturity, as do the absolute unconditional covariances of the excess returns with the macroeconomic variables..

Descriptive statistics for $\pi_{3,t+1}$, $\pi_{6,t+1}$ and $\pi_{12,t+1}$ are presented in Table 5.1 of this paper. We see that, though the mean inflation increases with the horizon, the standard deviation decreases. This is to be expected, since the longer the period over which we take the growth rate, the smoother the series will be. Table 5.2 reports the unconditional sample correlations of the series used in this paper. The unconditional correlations of the macroeconomic variables with the excess returns are negative and for each macroeconomic variable they increase in absolute value with the maturity of the bond. However, interestingly, the covariances of the excess return of any given maturity do not seem to have a monotonous relation with the horizon in the future over which realised inflation has been calculated. For example the unconditional covariance of the excess return on the 3-month bond with 3-month ahead inflation is -0.04, with 6-month ahead inflation is it -0.08

[7] Estimation results are available upon request.

while with 12-month ahead inflation it is -0.07. Even more important for our estimation will prove to be the fact that although the conditional covariance between the excess return on the 6-month bond and 6-month ahead inflation is very close to that between the excess return on the 12-month bond and 12-month ahead inflation, at about -0.15, the unconditional covariance between the excess return on the 3-month bond and 3-month ahead inflation is much lower, at about -0.04. This will be reflected in our estimation results.

5 Discussion of Results

All estimation results are reported in Tables 5.3 to 5.6. Table 5.3 reports the estimates of the Fisher equation of the estimation for the 1-year horizon. The excess holding-period returns included in this estimation are of bonds maturing in 1 and 5 years. However, in accordance with equation 20, it is only the covariances of the excess return of the 1-year maturity that have been included in the Fisher equation. Row 1 estimates are of the Fisher equation where no constant or risk premium term has been included. In other words, only the 1-year yield has been included in the conditional mean of 1-year ahead inflation. The coefficient of this term is 0.566, highly significant but significantly different from the expected value of 1. The fit of this equation is poor, though better than that of the corresponding univariate estimations usually performed (see for example x, y, z).

The second row reports the results of the estimation where the risk premium terms, have been included in the Fisher equation, i.e. the estimated conditional covariances of the 1-year excess holding-period return with each of the two macroeconomic variables. The estimated coefficient of the 1-year yield is 0.862, still highly significant. It still has the correct sign and, though statistically it remains significantly different from its theoretical value of unity, it is much closer to it now than it was before the inclusion of the risk premium term. Hence it seems that the inclusion of the risk premium proxy reduces the bias observed in the standard specification of the Fisher equation. The coefficients of the conditional covariances included in the mean are in the following two columns. The coefficient of the inflation covariance is estimated to be 0.76, and is highly significant. The coefficient of the consumption covariance is much smaller, at 0.1 approximately, but is also statistically significant. Furthermore, their inclusion is jointly highly significant, as demonstrated by the likelihood ratio test comparing the two specifications. Its value is 26.8, much higher than the critical value of 3.84 (for one constraint) at a 5% significance level. One might argue that these estimated coefficients have the expected sign; they are both positive which, given the negative unconditional correlation of the 1-year excess return with both macroeconomic variables, implies that the risk premium proxy enters the Fisher equation, on average, negatively, as one might expect based on equation 12. Both the R-squared and the share of variance explained are substantially higher than before, at 0.86 and 0.69 respectively[8] . Figure 5.1 shows the fit of this specification to inflation data.

Finally, row 3 reports the results of our null estimation where, as above, the risk premium terms have been included, and the coefficient of the yield has now been constrained to equal its theoretical value of unity. Our results are similar. The conditional covariance terms have, as before, positive estimated coefficients and remain highly significant. However both are now slightly higher, with the inflation covariance coefficient estimated at 0.845 and the consumption covariance coefficient having increased to 1.42. Furthermore, both the share of variance explained by this specification and the R-squared of the Fisher equation have increased to 0.9 and 0.73 respectively.

Figure 5.2 presents graphically the fit of all three estimations. Both specifications including a risk premium proxy offer a remarkably better fit than the first estimation, while the null estimation clearly provides the best fit.

An examination of the above results might lead us to conclude that the null hypothesis, i.e. the Fisher equation including the risk premium proxy and constraining the yield coefficient to unity,

[8] Since no constant has been included in our estimati, R-squared may be biased. Hence, we also report the share of inflation's variance which is explained by each specification.

is in fact the best specification for the inflation equation. However, the likelihood ratio criterion forces us to question this conclusion, as this specification is rejected against the previous, less restricted one. Hence, although the inclusion of the risk premium terms is clearly significant and substantially improves the predictive power of the Fisher equation, given the statistical evidence of both the likelihood ratio tests and the t-student tests, we cannot actually constrain the yield coefficient to its theoretical value. This could be because a relevant variable has been omitted from the specification, i.e. the real rate, or because the risk premium term is a simply a proxy of the rolling risk premium and does not capture its every fluctuation. However, one can conclude that the Fisher equation provides a much better predictor of future inflation when the risk premium is taken into account, by including an appropriate risk premium proxy.

Figure 5.3 shows the contribution of the risk premium terms for each of the two specifications, for the 1-year horizon. As expected, the contribution is mostly negative in both estimations. Nevertheless, in both specifications, the evolution of the risk premium through time differs substantially from that estimated in Balfoussia and Wickens (04) and, at first glance, seems hard to interpret. From the late 70's onwards, it seems similar to what we might expect, given the risk premia estimated in Balfoussia and Wickens (04). However, during the seventies, we now see two peaks, one just before and during the first part of the Fed's "monetary experiment" and a second one earlier, around the time of the first oil crisis. One might speculate that this could be attributed to the risk premium term also picking up part of the omitted real rate. If, in fact, the underlying unobservable real rate is not constant but time varying (as is highly likely over the first half of our sample) then it may be that the risk premium term is proxying not only for the risk premium contribution but for the risk premium adjusted by the real rate. Indeed, around the time where we see the peaks, the *ex post* real 1-year rate was actually negative, as a result of the excessively high inflation rates. Hence, its contribution to the Fisher equation would be positive, rather than negative, as one would expect under normal economic circumstances. One might conclude that, though during the 80's and 90's the real rate could be taken to be roughly constant, this may not be a reasonable assumption for the decade of the 70's, as a result of both the oil crisis and the Fed's sudden change of monetary policy towards the end of the decade. This might explain the rather unexpected shape of the contribution of the covariance terms to the Fisher equation. If this is indeed the case, then the rejection of the $\delta_{n,1} = 1$ hypothesis by t-tests and the likelihood ratio test may in fact be picking up this misspecification, rather than actually rejecting the null hypothesis of the Fisher equation.

Tables 5.4 and 5.5 show the estimates of corresponding estimations for the 6-month and 3-month horizons. The estimates and conclusions are broadly similar. Thought we cannot accept the hypothesis that the coefficient of the yield equals unity, the Fisher equation consistently performs much better at all horizons, once a proxy for the time-varying rolling risk premium has been included in the estimation. The only slight deviation is the very short end of the spectrum, at the 3-month horizon, where the fit, though clearly improved, is not as good, and the coefficient of the yield is 0.67, not as close to unity as at longer maturities. One possible interpretation of this different behaviour may be that the 3-month bond is very close to a risk-free asset. Hence, it may be traded differently than longer maturity bonds, typically much more heavily. Furthermore, at such short horizons there is likely to be much more noise than at longer ones, possibly making it harder to extract the inflation component from the yield. Nevertheless, the risk premium elements are highly significant, as for other horizons, and the specification is clearly superior to that which does not include a proxy for the rolling risk premium. Figures 5.3, 5.4, 5.5, 5.6 present the fit and risk premium contribution of the estimations with a 6-month horizon and a 3-month horizon. Once again the conclusions drawn are along the same lines. It is worth noting the increased volatility of the 3-month risk premium contribution, reinforcing our conclusion that the poorer fit at this maturity is due to the presence of noise.

Table 5.6 presents one representative full set of estimated parameters, that of the estimation for the 6-month horizon[9] . We see that the coefficients of the inflation covariance in the excess

[9] The complete estimates for all specifications are available upon request.

return equations are highly significant and similar to those estimated in Balfoussia and Wickens (04), at -0.519 for the 6-month maturity and -0.531 for the 2-year maturity. Interestingly, this time the corresponding coefficients for TPC are also negative, though not significantly so, as would be expected by the theory. Figure 7 presents the holding-period risk premia, that is the risk premia included in the excess holding-period return equations for the maturities included in this horizon. Despite the slightly modified SDF, these are broadly similar to the ones obtained with an AR(1) representation of inflation, if slightly higher towards the end of the sample. The constant and first lag of TPC in the corresponding equation are highly significant. Turning to the ARCH and GARCH estimates, we see they are also significant, for the excess holding-period return equations as well as for the macroeconomic variables, often with the t-statistics of two or more significant places. Concerning the dynamic structure of the conditional variance-covariance matrix, the estimated GARCH parameters are much larger than the ARCH parameters, indicating that the conditional covariance structure of excess holding-period returns depends largely on the lagged conditional covariance matrix, and much less on lagged innovations. Typically estimates are close to 0.9 and 0.3 respectively.

6 Conclusions

In this paper we have tried to integrate the workhorse of inflation prediction, the Fisher equation, with the stochastic discount factor theory. This allows us to construct a multivariate framework in which to jointly model the term structure and inflation, while generating a suitable risk premium proxy for each inflation prediction horizon. The inclusion of the rolling risk premium proxy seems to be highly significant. It substantially improves the predictive power of the Fisher equation at all horizons, while clearly reducing the bias of the estimated yield coefficient. We thus provide strong evidence in support of the Fisher equation and we conclude that it provides a sound theory and a useful modelling tool for inflation, once the risk component has been appropriately taken into account.

7 References

Fisher I. 1930. The Theory of Interest. New York: A. M. Kelly.

Table 5.1: Descriptive Statistics

		Mean	Median	Maximum	Minimum	Std. Dev.	Skewness	Kurtosis
Growth rates	CPI3	5.13	4.09	17.73	-3.65	3.51	0.97	0.69
	CPI6	5.19	4.29	15.65	-0.18	3.28	1.09	0.59
	CPI12	5.31	4.25	14.76	1.10	3.19	1.14	0.42

Notes
1. All series are in annualised percentages.
2. CPI6 is the 6-month ahead, annualised *ex-post* change in CPI inflation, CPI12 the same for a 12-month horizon etc.

Table 5.2: Sample Correlations

		3 months	6 months	1 year	2 years	5 years	CPI3	CPI6	CPI12	TPC
Excess holding period returns	3 months	1.00								
	6 months	0.87	1.00							
	1 year	0.73	0.95	1.00						
	2 years	0.64	0.88	0.96	1.00					
	5 years	0.54	0.76	0.87	0.94	1.00				
Growth rates	CPI3	-0.04	-0.12	-0.17	-0.21	-0.24	1.00			
	CPI6	-0.08	-0.15	-0.19	-0.23	-0.25	0.92	1.00		
	CPI12	-0.07	-0.12	-0.15	-0.18	-0.20	0.85	0.94	1.00	
	TPC	-0.02	-0.05	-0.07	-0.07	-0.09	-0.16	-0.12	-0.11	1.00

Table 5.3: Fisher equation estimation results Inflation prediction horizon: 1 year							
Constraints imposed	R_{12}	CPI12	TPC	Likelihood ratio tests[1]		Share of variance[2]	R-squared
				No risk premium	Null		
No risk premium term	0.566 * *50.97*	-	-			0.23	0.52
Incl. risk premium term	0.862 * *19.64*	0.760 *15.01*	0.105 *4.28*	26.8 #		0.86	0.69
Null	1.000 .	0.845 *25.78*	0.142 *6.50*		133.5 #	0.90	0.73

Notes

1. Likelihood ratio tests:
   The column title corresponds to the restriction tested.
   # denotes rejection of the restriction tested, at the 5% significance level.
2. Share of inflation variance explained in each estimation
3. *t*-statistics are below the estimated parameters in italics.
4. * denotes rejection of the hypothesis of coefficient equality to unity, using a *t*-test at the 5% significance level.

Table 5.4: Fisher equation estimation results Inflation prediction horizon: 6 months							
Constraints imposed	R_6	CPI6	TPC	Likelihood ratio tests[1]		Share of variance[2]	R-squared
				No risk premium	Null		
No risk premium term	0.552 * *42.51*	-	-			0.20	0.45
Incl. risk premium term	0.985 * *21.02*	1.433 *9.23*	0.069 *2.07*	49.8 #		0.61	0.67
Null	1.000 .	1.613 *15.66*	-0.018 *-1.09*		7.2 #	0.66	0.68

Notes

1. Likelihood ratio tests:
   The column title corresponds to the restriction tested.
   # denotes rejection of the restriction tested, at the 5% significance level.
2. Share of inflation variance explained in each estimation
3. *t*-statistics are below the estimated parameters in italics.
4. * denotes rejection of the hypothesis of coefficient equality to unity, using a *t*-test at the 5% significance level.

Table 5.5: Fisher equation estimation results Inflation prediction horizon: 3 months							
Constraints imposed	R_3	CPI3	TPC	Likelihood ratio tests[1]		Share of variance[2]	R-squared
				No risk premium	Null		
No risk premium term	0.590 *	-	-			0.19	0.46
	40.42						
Incl. risk premium term	0.670 *	-0.279	0.293	102.1 #		0.28	0.57
	28.60	*-3.11*	*1.63*				
Null	1.000	0.915	4.341		6.7 #	0.60	0.80
	.	*7.94*	*11.09*				

Notes

1. Likelihood ratio tests:
   The column title corresponds to the restriction tested.
   # denotes rejection of the restriction tested, at the 5% significance level.
2. Share of inflation variance explained in each estimation
3. *t*-statistics are below the estimated parameters in italics.
4. * denotes rejection of the hypothesis of coefficient equality to unity, using a *t*-test at the 5% significance level.

Table 5.6: Complete Estimation Results Bond maturities: 6 months & 2 years Inflation prediction horizon: 6 months Risk premium term included							
Conditional mean equations							
Equation		Constant	Own lag	Conditional Covariances[2]			
				6 months	2 years	CPI6	TPC
Excess returns	6 months	0.000	0.000	-	-	-0.519	-0.055
		.	.	.	.	*-6.30*	*-1.54*
	2 years	0.000	0.000	-	-	-0.531	-0.007
		.	.	.	.	*-6.75*	*-0.16*
Macro variables	CPI6	0.000	0.985	1.433	-	-	0.069
		.	*21.02*	*9.23*	.	.	*2.07*
	TPC	4.342	-0.261	-	-	-	-
		9.16	*-5.17*	.	.	.	.
Conditional variance equations							
Equation		Long run variance-covariance matrix[3]				ARCH	GARCH
Excess returns	6 months	10.315	35.787	-1.151	-1.854	-0.320	0.927
		.	.	.	.	*-18.04*	*114.10*
	2 years	35.787	160.706	-6.760	-8.731	-0.250	0.935
		.	.	.	.	*-11.62*	*87.80*
Macro variables	CPI6	-1.151	-6.760	8.525	-1.210	0.175	0.981
		.	.	.	.	*7.69*	*169.59*
	TPC	-1.854	-8.731	-1.210	49.516	0.292	0.710
		.	.	.	.	*5.86*	*6.57*

Notes

1. *t*-statistics are below the estimated parameters in italics.
2. The coefficients are of the conditional covariances between the variables defined by the column and row of each cell.
   An exception is the second in-mean covariance included in the Fisher equation. The coefficient reported is not that of the covariance between the two macroeconomic variables but that between TPC and the 6-month excess holding period return.
3. A consistent estimator of the long-run variance covariance matrix, to which H_0 is subsequently fixed, is obtained by estimating a standard homoskedastic VAR estimator for the whole system. Hence no t-statistics are reported for H_0.

Figure 5.1: The Fisher Equation Fit to 12-month Ahead Inflation

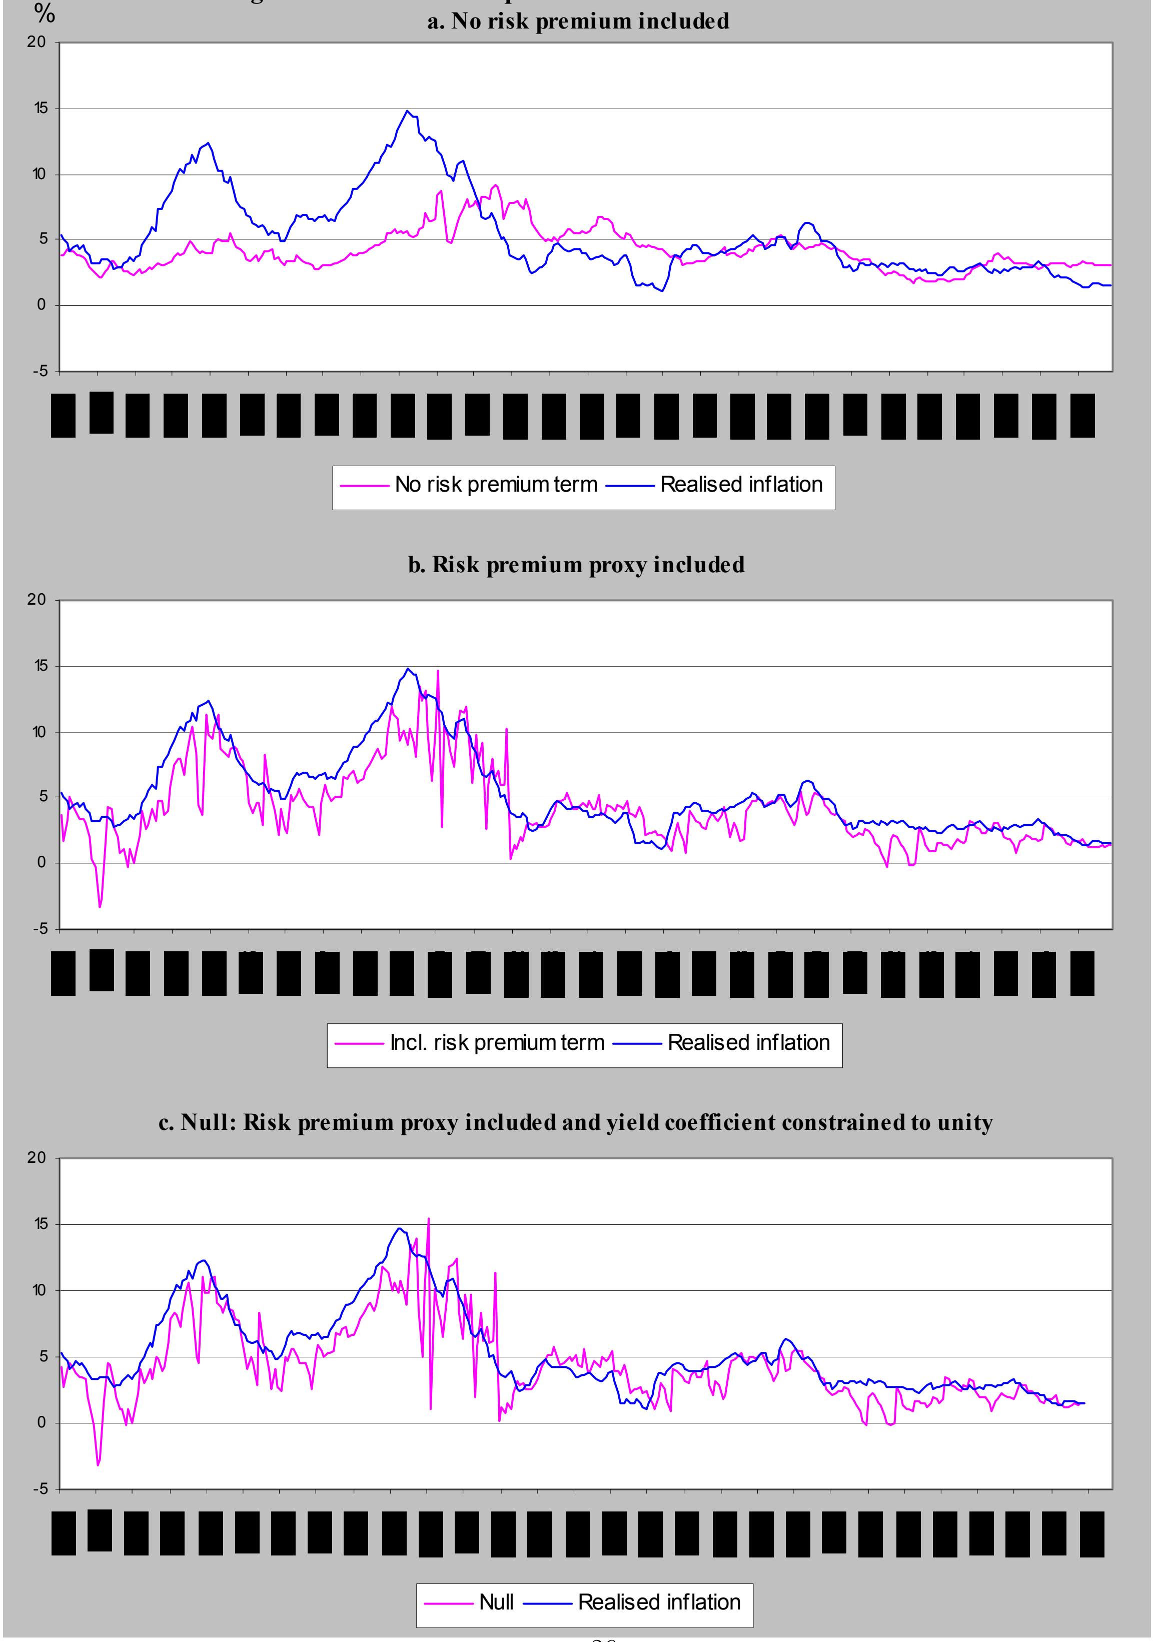

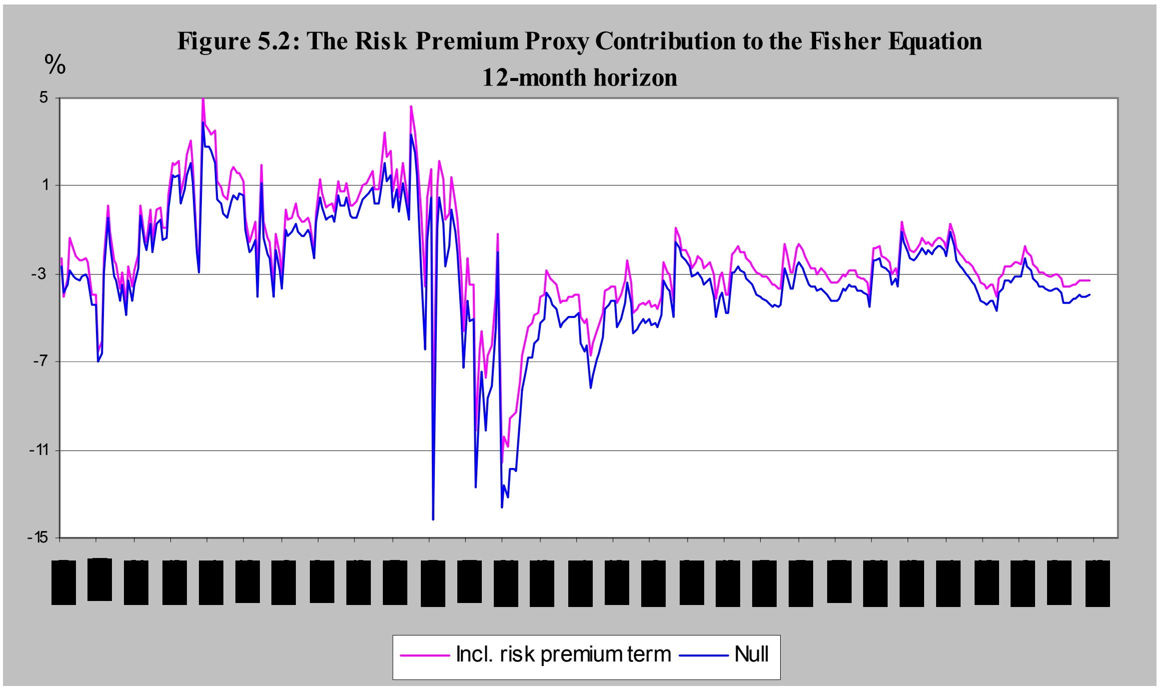

Figure 5.2: The Risk Premium Proxy Contribution to the Fisher Equation
12-month horizon

Figure 5.3: The Fisher Equation Fit to 6-month Ahead Inflation

a. No risk premium included

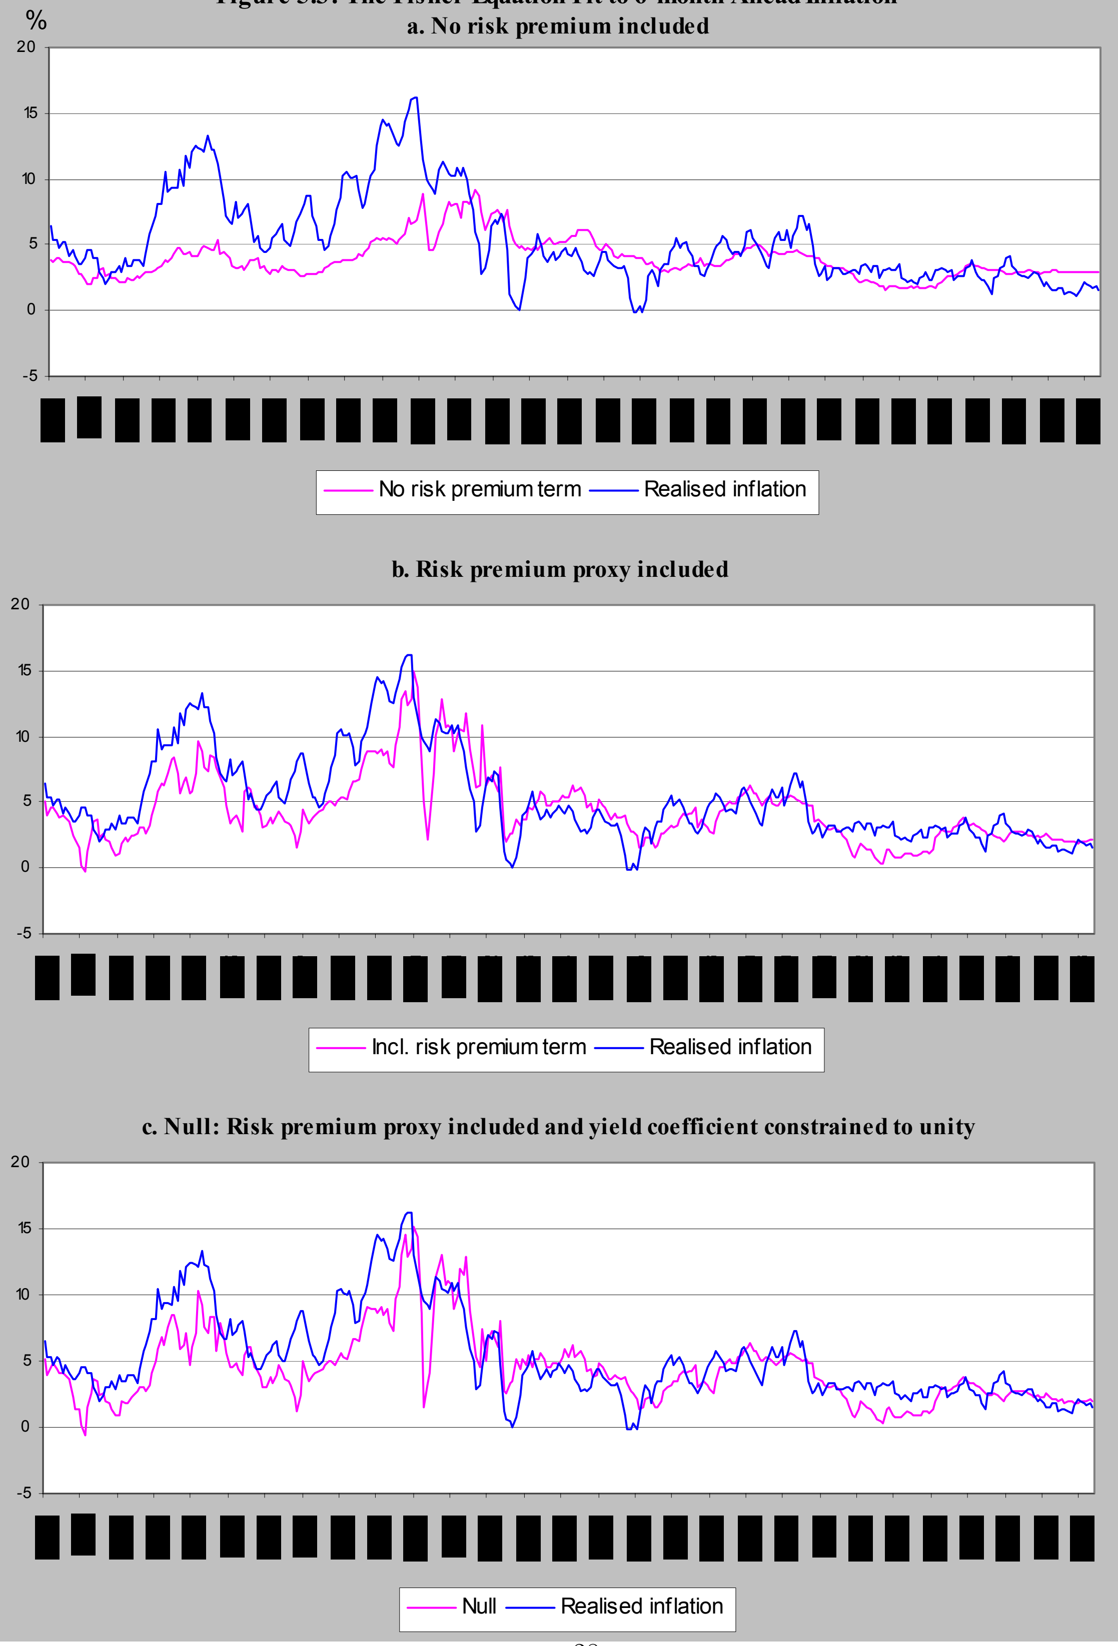

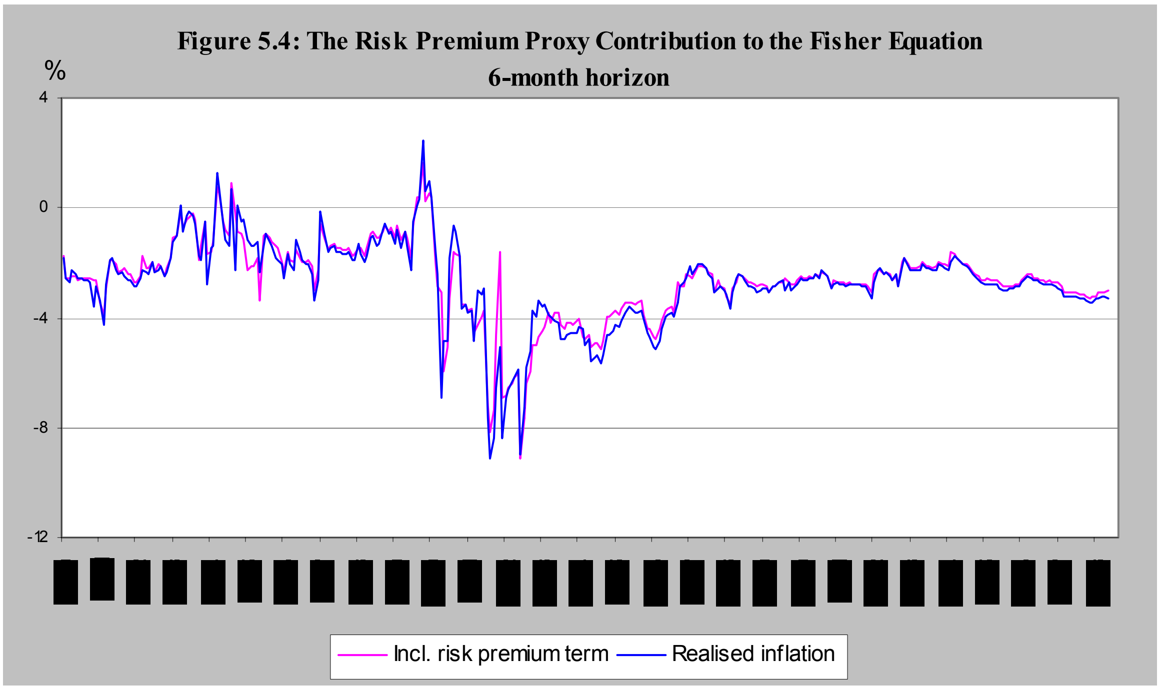

Figure 5.4: The Risk Premium Proxy Contribution to the Fisher Equation
6-month horizon

Figure 5.5: The Fisher Equation Fit to 3-month Ahead Inflation

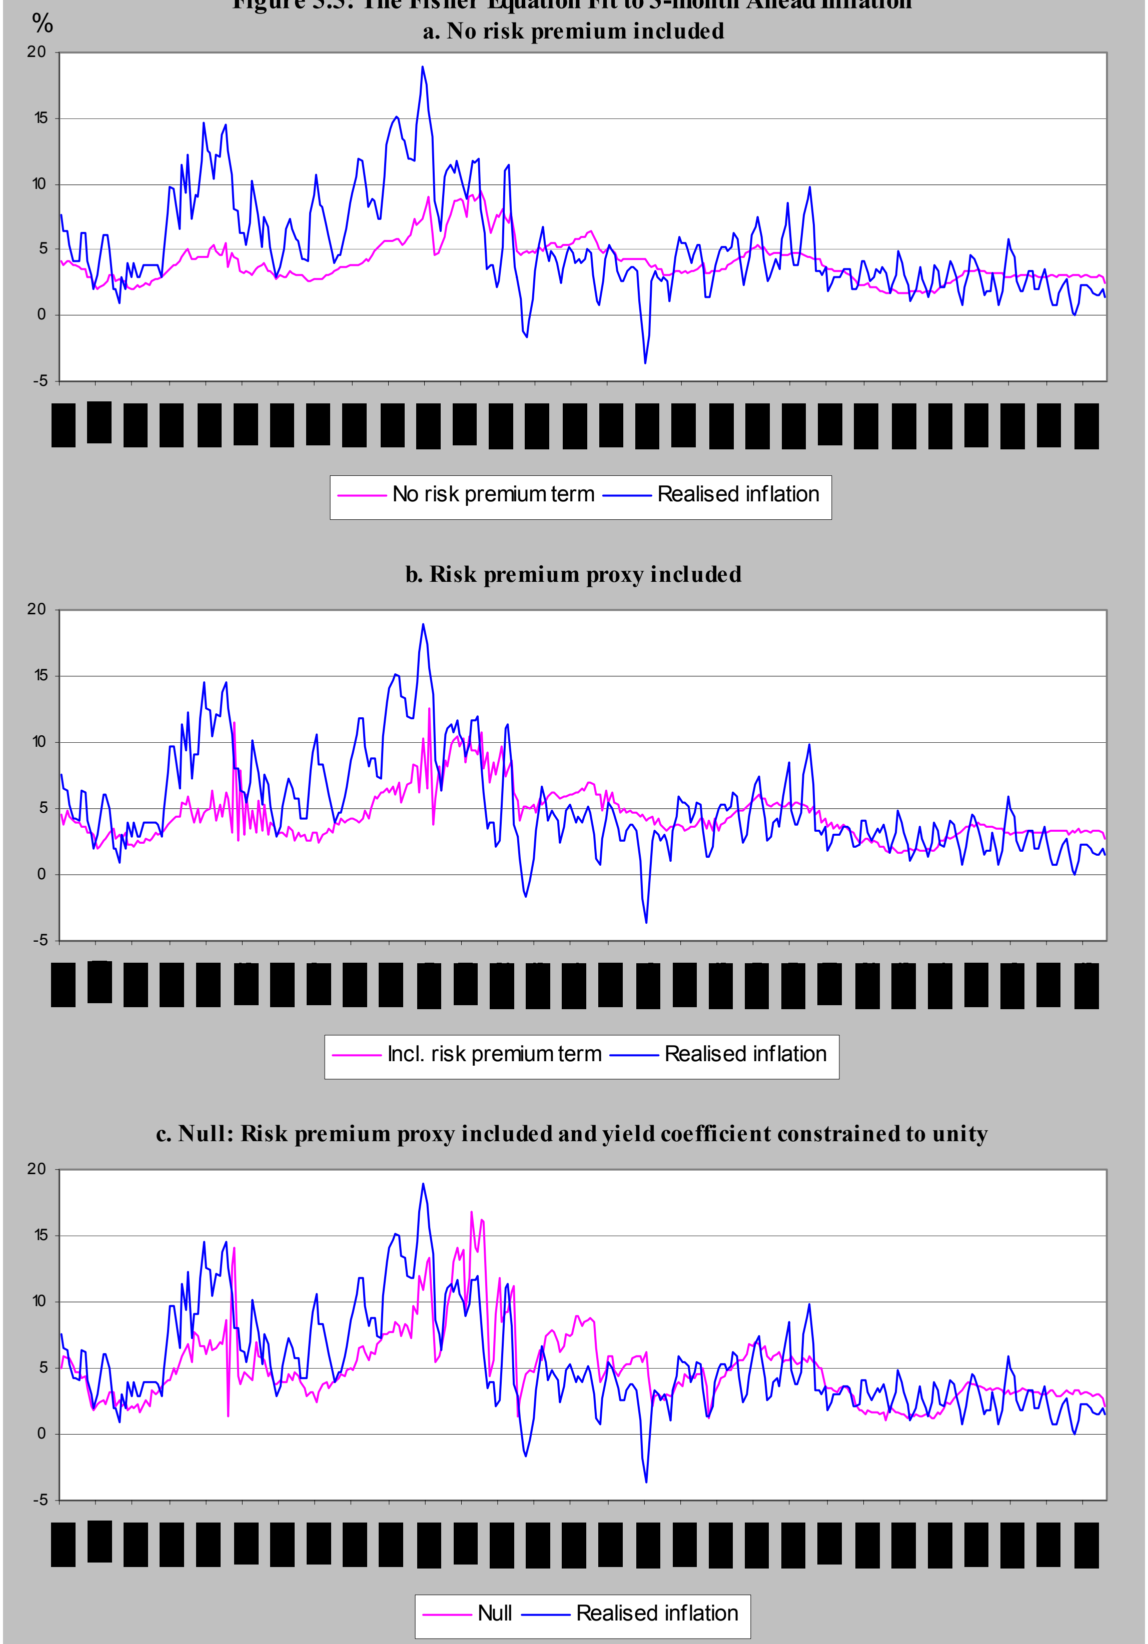

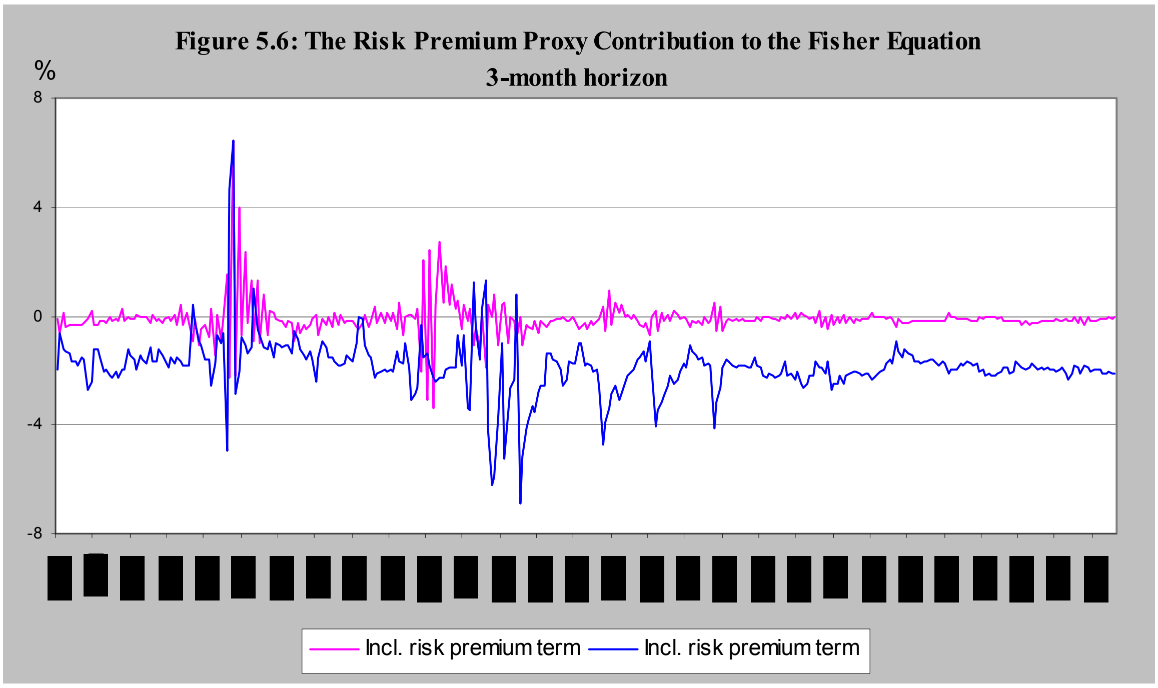

Figure 5.6: The Risk Premium Proxy Contribution to the Fisher Equation
3-month horizon

Figure 5.7: Estimated Risk Premia and Excess Returns
6-month Inflation Prediction Horizon - Risk Premium Proxy Included