Expected Inflation Risk Premium, News Decompositions and the Cross Section of Asset Returns

Alexandru Rosoiu [1]

Fuqua School of Business

Duke University

November 20, 2015

[1]E-mail: alexandru.rosoiu@duke.edu

Abstract

I show that expected inflation risk is priced in the cross section of stock returns even after controlling for cash flow growth and volatility risks. Motivated by this evidence I study a long run risk model with a built-in inflation non-neutrality channel that allows me to decompose the real SDF into news about current and expected cash flow growth, news about expected inflation and news about volatility. The model can successfully price a broad menu of assets and provides a setting for analyzing cross sectional variation in expected inflation risk premia. For industries like retail and durable goods inflation risk can account for nearly a third of the overall risk premium while the energy industry and a broad commodity index act like inflation hedges. Nominal bonds are exposed to expected inflation risk and have inflation premiums that increase with bond maturity. The price of expected inflation risk was very high during the 70's and 80's, but has come down a lot since being very close to zero over the past decade. On average, the expected inflation price of risk is negative, consistent with the view that periods of high inflation represent a "bad" state of the world and are associated with low economic growth and poor stock market performance.

1 Introduction

How important is inflation risk for the economy and why should one care about it? One potential way of answering this question would be to look at the welfare cost of inflation and the large macroeconomic literature covering the topic. Early papers like Fisher (1981) and Lucas (1981) follow the tradition of Bailey (1956) and Friedman (1969) and view real money balances as a consumption good and inflation as a tax on real balances. The welfare cost implied by this approach is small. An increase in inflation from 0% to 10% would have a cost of only 0.30% - 0.45% of GNP depending on the measure of money used, an estimate revised upwards to a little under 1% by Lucas (2000). Cooley and Hansen (1989) argue that if labor is supplied elastically, an inflation tax can cause agents to substitute from activities that require cash (like consumption) into activities that do not (like leisure) leading to an inefficient allocation of resources. Lucas (2000) argues that the welfare effect of this decision is small. One crucial assumption in his setting is that the steady state growth rate of output is independent from monetary policy. Dotsey and Ireland (1996) relax this assumption allowing the allocative effects of inflation to affect the equilibrium growth rate as well as the level of aggregate output. Although the effect on the growth rate is small, the welfare implications are large nearly doubling the estimates of Lucas (2000). Finally, Lagos and Wright (2005), using a search framework, argue the welfare cost of inflation is much higher in the region of 3%-4%.

Regardless of which estimate one perceives as more accurate, it is clear that inflation has strong implications for welfare and that economic agents strongly dislike states of high inflation. One would then expect to see this behaviour reflected in the risk premia across various asset classes. Specifically, assets that do badly in times of high inflation should be perceived as riskier by investors and should command a premium. It is precisely this observation that motivates the current paper.

I begin by showing that expected inflation risk is indeed priced in the cross section of US stock returns even when controlling for well established sources of risk like cash flow growth and volatility. I achieve this by first estimating the conditional betas at security level using a rolling window time series regression. I then perform a three dimensional conditional sort into 27 bins (3 x 3 x 3) and construct 1-quarter ahead value-weighted portfolio returns for each of them. I collapse the whole structure by averaging across two of the dimensions and reach three final portfolios that give different exposure to the chosen risk source while controlling for the others. A long-short strategy over the sample period 1974 - 2014 achieves a spread of about 2% annualized for expected inflation risk and of 3.6% and 3.7% respectively for cash flow growth and volatility risks. Focusing on the long-short portfolio giving exposure to expected inflation risk I find a spread as high as 5% for the period 1974-1987 when inflation expectations were high. This spread drops to about 2% over the next decade and becomes negative but very small over the past 15 years when inflation has become procyclical.

This exercise bears some resemblance to Duarte (2013). He constructs an expanding window, Vasicek adjusted inflation beta for each security in the CRSP universe[1] where the inflation shock is measured as the first difference of CPI inflation. He performs a double sort on size and the inflation beta and then collapses the size dimension. His resulting portfolios exhibit a spread in inflation risk while having roughly equal exposure to the Fama French risk factors. My paper on the other hand aims at controlling for other sources of macroeconomic risk. Moreover I use a different methodology, set of variables and I aim to match a broader cross section of assets with the model.

Weber (2015) matches a BLS dataset (underlying the PPI) to CRSP and Compustat

[1] As opposed to the more commonly used NYSE, NASDAQ and AMEX securities with common share codes.

firms and sorts on the frequency of price adjustments. He finds that firms that adjust their prices infrequently earn a return premium of 2%-4% a year compared to firms that are more flexible. He shows the return differential is successfully explained by a cash flow beta with stickier firms being more exposed to cash flow risk and having more volatile cash flows in general. Compared to him I sort on expected inflation rather than price stickiness and I argue that expected inflation is priced in the cross section even after controlling for cash flow and volatility risks. I interpret this as evidence of expected inflation risk directly affecting the marginal utility of the agent and I study a model where this risk source shows up in the real stochastic discount factor for the economy.

Consistent with the results from the portfolio sorting exercise, in the data, periods of high expected inflation are associated with low expected economic growth. Figure 1 plots a 12-quarter trailing mean for the personal consumption expenditure (PCE) annualized real growth rate along with that for the associated PCE deflator inflation over the period 1952 - 2014. The countercyclical relationship between these variables is striking and spans over 50 years of data until the early 2000's when inflation seems to have become procyclical. It is worth mentioning that this behaviour is robust to the choice of variables and one would get an almost identical plot by replacing PCE real growth by real consumption growth (excluding durable goods) or by redoing it altogether in terms of real GDP growth and change in the GDP price index.

[**Place Figure 1 about here**]

I continue by studying a long run risk model that can reconcile the above facts. Having a model provides theoretical motivation for the factors I consider while, at the same time, pinning down their prices of risk in terms of the agent's preference parameters and providing testable implications. The mechanism that enables expected inflation to directly affect the agent's marginal utility and to be a priced source of risk is an inflation non-neutrality channel (expected inflation forecasts expected consumption growth with a negative coefficient). This modeling device has been used before in both Piazzesi and Schneider (2007) and Bansal and Shaliastovich (2013). Piazzesi and Schneider (2007) aim at matching the term structure of nominal bonds. Bansal and Shaliastovich (2013) include it in a long run risk model with time-varying volatilities of expected growth and expected inflation. They show such a model can simultaneously account for bond return predictability as well as uncovered interest rate parity violations. Compared to their paper I entertain a single volatility source and I take the model in a different direction both in the manner I take it to the data and through the set of facts I am aiming to explain. Finally, Kung (2015) sets up an economy where the production and price-setting decisions of firms drive low-frequency movements in growth and inflation rates that are negatively related, endogenizing the non-neutrality channel I use.

I decompose the innovation to the real SDF into news about current and expected consumption growth, news about expected inflation and news about volatility and pin down the news components in the data using a standard VAR approach. This methodology is closest to Bansal et al. (2014) and Campbell et al. (2015). Proxying expected inflation by the SPF survey inflation forecast allows me to identify the expected inflation shock and, by imposing the model restrictions, to disentangle the news about expected inflation and news about expected cash flow components. With reasonable preference parameters estimates (a relative risk aversion coefficient of $\gamma = 4.13$ and intertemporal elasticity of substitution of $\psi = 2.2$) the model does a good job at matching the excess returns on a wide cross section of assets including industry portfolios, nominal bonds and a commodity index. For industries like retail or durable goods the inflation premium is in excess of 2.5% and represents nearly a third of the overall premium on this portfolios. At the opposite end the portfolio formed of energy (mostly oil) companies acts like an inflation hedge and has a small and negative inflation risk premium. The model identifies an inflation risk premium in nominal bonds that goes up with bond maturity and finds the commodity index acts as an inflation hedge. It does well also in term of matching size and book to market portfolios, although those turn out to be not particularly interesting from an inflation perspective. The price of expected inflation risk has changed considerably during my sample period being very high during the 70's and 80's and coming down a lot since and hovering close to zero over the past decade.

My paper fits into a larger stream of literature investigating the relationship between stock returns and inflation. Classical economic intuition holds that stocks are claims to physical capital and therefore they should be an inflation hedge in the sense that real return on stocks should be independent from (or at least uncorrelated with) the rate of inflation. This view implies the beta in the univariate regression of real stock returns on inflation should be 0. Several papers point out that this basic prediction does not hold in the data. For example, Bodie (1976) finds that, for the sample period 1953 to 1972, "contrary to a commonly held belief among economists, the real return on equity is negatively related to both anticipated and unanticipated inflation" and therefore in order "to use common stocks as a hedge against inflation one must sell them short." Fama and Schwert (1977) investigate the hedging inflation properties of various asset classes and reach the same conclusion as far as common stock returns are concerned. Schwert (1981) analyzes the reaction of daily returns on the S&P portfolio to CPI announcements for the period 1953 - 1978. He finds that aggregate stock returns react negatively to unanticipated inflation, but the effect is dispersed across the 15 trading days surrounding the announcement rather than concentrated on the day of the BLS release. More recently Bekaert and Wang (2010) look at the inflation hedging properties of 48 country equity indexes and find that most of them are exposed to inflation risk. While emerging market equity returns do somewhat better than developed markets, they still provide a poor overall inflation hedge.

This evidence of stock market underperforming in times of high realized and expected inflation is further corroborated by the findings of Bekaert and Engstrom (2010). Their paper argues that the high correlation between bond yields and equity yields is generated by the inflation risk premium and the stock risk premium being strongly positively correlated. Among other things they show that if one decomposes the equity yield into expected cash flow growth, expected future risk free rates and expected future risk premium the correlation between expected inflation and the last term (interpreted as long-run risk premium) is strikingly large. Bekaert and Engstrom (2010) also show that the correlation between the inflation and the stock risk premiums is related to the incidence of stagflation (countries with more frequent stagflation episodes display a stronger positive correlation between the two premia). Campbell, Sunderam, and Viceira (2013) discuss the changing correlation between the returns on stocks and nominal bonds. They attribute this to the changing correlation of inflation with the real economy. Compared to these papers I take things a step further by allowing expected inflation risk to directly show up in the real SDF of the representative agent. This generates an expected inflation risk premium for all assets priced in the economy and allows me to both quantify the premium as well as look at its cross sectional and time series variation.

The remaining of the paper is organized as follows. Section 2 presents evidence that expected inflation risk is priced in the cross section of US stocks even after controlling for cash flow growth and volatility risks. In section 3 I study a long run risk model with a built in inflation non-neutrality channel. The model provides theoretical motivation for the factors considered and pins down the market prices of risk for each of them. Section 4 shows how to use standard VAR methods in order to pin down the parameters of the reduced form model and construct from the data the news terms that drive the SDF. I then discuss the empirical implications of the model. Section 5 concludes.

2 Evidence from Portfolio Sorts

In this section I show that inflation risk is priced in the cross section of US stock returns even after controlling for established sources of risk like cash flow growth and market volatility. Specifically, I will sort US stocks into bins based on their exposure to cash flow growth risk, expected inflation risk and volatility risk and I will form portfolios that give exposure to one risk source while controlling for the other two. The premiums on these portfolios are consistent with the idea that low cash flow growth, high expected inflation and high realized volatility are viewed by investors as the bad states of the world. Put differently, cash flow growth has a positive price of risk while expected inflation and volatility have negative prices of risk when the full sample is considered.

2.1 Data

Expected inflation is taken from the Philadelphia Fed's Survey of Professional Forecasters. I use the median forecasts for the GDP price index because it is available for a longer horizon (one quarter ahead forecasts are available going back to 1968 Q4). The survey comes out quarterly and as a result this is the frequency chosen for the exercise in this section. Cash flow growth is measured as real dividend growth for the CRSP value weighted market portfolio

(NYSE, NASDAQ and AMEX)[2]. I deflate the quarterly dividend growth using GDP price index inflation taken from BEA website. I use the same realized market volatility measure as Bansal et al. (2014), constructed by summing up the squared monthly log real returns on the market portfolio during each quarter[3]. I use the 3-month Treasury bill rate taken from the FRED database as my measure of quarterly risk free rate. Finally, individual security returns are taken from the CRSP database.

2.2 Conditional Sort

I begin by measuring the exposure of each security in my sample to the three sources of risk at each point in time. I do this by running rolling window time series regressions in the spirit of Fama and MacBeth (1973) and Black, Jensen, and Scholes (1972). Sitting at time t, I only use information up to time t-1 to account for the fact that inflation for the current quarter t wouldn't be known to the investor in real time and therefore the real dividend growth and volatility measures described above might not be available. In order to be included, a security must (as it is typical in the literature) be listed on NYSE, NASDAQ or AMEX and have a common share code (10 or 11), valid price and number of shares outstanding at time t. I also require securities to have no missing returns during the past 20 quarters. Finally, I exclude all securities that delist at the end of the current quarter.

The procedure consists of three steps. First, I run a small VAR on the sample period t-21 to t-1 to extract the cash flow, expected inflation and volatility shocks over a 20 quarters period. The VAR variables consist of real dividend growth Δd_t, log price dividend ratio pd_t, realized market volatility RV_t, one quarter ahead expected inflation $x_{\pi,t} = E_t\pi_{t+1}$, and two

[2]CRSP provides data on the return with and without dividends for a broad market portfolio. I use these returns to infer monthly level dividends which I add up during each quarter. Finally a 4 quarter trailing mean is taken to deseasonalize as it is typical in the literature.

[3]I use CPI inflation taken from the BLS website to deflate the monthly market returns

quarter ahead expected inflation $E_t x_{\pi,t+1} = E_t(E_{t+1}\pi_{t+2}) = E_t\pi_{t+2}$. I impose two type of restrictions on the VAR. On one hand I read the shock to expected inflation as the forecast update. On the other I never include both one quarter and two quarters ahead inflation forecasts on the right hand side of the same regression to avoid a collinearity problem.

$$\begin{pmatrix} \Delta d_{t+1} \\ x_{\pi,t+1} \\ RV_{t+1} \\ pd_{t+1} \\ E_{t+1}x_{\pi,t+2} \end{pmatrix} = \begin{pmatrix} * & * & * & * & 0 \\ 0 & 0 & 0 & 0 & 1 \\ * & * & * & * & 0 \\ * & * & * & * & 0 \\ * & 0 & * & * & * \end{pmatrix} \begin{pmatrix} \Delta d_t \\ x_{\pi,t} \\ RV_t \\ pd_t \\ E_t x_{\pi,t+1} \end{pmatrix} + \begin{pmatrix} \epsilon_{\Delta d_{t+1}} \\ \epsilon_{x_{\pi,t+1}} \\ \epsilon_{RV_{t+1}} \\ \epsilon_{pd_{t+1}} \\ \epsilon_{E_{t+1}x_{\pi,t+2}} \end{pmatrix} \tag{1}$$

Second, for each security i I regress the security excess return on the VAR variables[4] and extract the firm level shocks over the same period considered above. Third and last, I run a contemporaneous regression of the firm level shocks $\epsilon_{R_{i,t}}$ during the window t-20 to t-1 on the cash flow growth, expected inflation and realized volatility shocks.

$$\epsilon_{R_{i,t}} = \beta_{i,t,\Delta d}\epsilon_{\Delta d_t} + \beta_{i,t,x_\pi}\epsilon_{x_{\pi,t}} + \beta_{i,t,RV}\epsilon_{RV_t} \tag{2}$$

Some of the betas estimated by this procedure are quite extreme. As a result I drop any security-date combination for which at least one of the betas estimated falls in the top 1% or bottom 1% of that beta's distribution. On average, due to overlap in extreme betas, about 5% of sample is trimmed at each point in time.

Given the conditional betas estimated above, I perform a triple conditional sort at each point in time. On each dimension I sort securities into three bins based on their exposure to that risk source: the "low" bin (L) containing securities with betas below the 3rd decile,

[4] I again exclude the last variable to avoid a collinearity issue.

the "medium" bin (M) containing securities with betas between the 3rd and 7th deciles and the "high" bin (H) containing securities with betas in excess of the 7th decile. I base the deciles only on stocks listed on the NYSE[5] as smaller securities tend to have more extreme betas and would dominate the low and high bins once NASDAQ securities become available in CRSP during the late 70s. More precisely, I first sort the securities into 3 bins based on their cash flow growth beta. Within each of the cash flow growth bins I then sort the securities contained in that bin into 3 additional bins based on their exposure to volatility risk. Finally, within each of the 9 existing bins (3 for cash flow growth $\times$ 3 for volatility) I sort the securities into 3 more bins based on their expected inflation betas. This procedures ensures that each security will end up in exactly one of the 27 final bins. The exposure of securities to each of the three risk sources will be similar for securities contained within the same bin, but will be different across bins. The order of the conditional sort (sorting on the inflation dimension last) is meant to provide a harder test for the existence of an expected inflation premium as one first controls for cash flow growth and volatility risk.

2.3 Portfolios

For each of the bins I now form a 1-quarter ahead value weighted portfolio (using delisting returns when available and necessary). This is consistent with a buy and sell strategy where at the end of a given quarter you estimate the security level conditional betas using available historical information, sort the securities in bins and then buy the securities in each bin in proportion with their market value. At the end of the next quarter you close your position and redo the exercise.

Ideally I would like to have portfolios that give me different exposure to a single risk

[5]Similar to the way Fama and French (1992) construct their size deciles.

source while controlling for the other two. Take for example expected inflation risk: one has 9 portfolios with "low" exposure to expected inflation, but different exposures to the other two risk sources. I combine these portfolios into a new equal weighted portfolio. I repeat the procedure for the 9 portfolios with "medium" and 9 portfolios with "high" expected inflation exposure. I now have only 3 portfolios that differ in their exposure to inflation risk, but have roughly equal exposure to the remaining risk sources as can be seen in figure 2.

[Place Figure 2 about here]

[Place Table 1 about here]

Table 1 presents summary statistics for these portfolios both for the full sample and across periods that are interesting from an economic point of view. In all cases the constructed portfolios exhibit large variation in terms of their average betas with respect to expected inflation risk, but have roughly similar exposures with respect to cash flow growth and volatility risk. The average cash flow growth beta is always positive and the average volatility beta is always negative consistent with typical economic intuition and other findings in the literature. Firms with a low (negative) exposure to expected inflation risk have a higher expected return compared with firms that have high (positive) loadings. Put differently, firms that do badly when an expected inflation shock hits the economy are viewed as riskier and investors require a premium for holding such companies. This return differential is approximately 2% on an annualized basis in the full sample which is quite sizeable and about a quarter of the 7.83% market risk premium over the same period. Between 1974 and 1987 when the US economy was experiencing a large overall level of inflation and frequent inflation shocks often culminating in recessions, a long-short strategy giving exposure to expected inflation risk would have netted investors as much as 5.42%. With inflation slowly being tamed, the return on the long-short portfolio comes down to 2.31% over the next 10 years (1988 - 1997). These findings are consistent with large expected inflation being the bad state of the world and expected inflation having a negative price of risk both in the full sample and across the first two subperiods considered. Finally, consistent with the idea that inflation has become pro-cyclical since in the late 90's and the results of Campbell, Sunderam, and Viceira (2013), I find that the return on the long-short expected inflation portfolio flips sign and gives a negative return of -0.73% from 1998 onward.

A similar exercise can be performed with respect to the other two sources of risk. One could first collapse the 27 portfolios into only 3 that give different exposure to cash flow growth risk while controlling for expected inflation and volatility exposure and then repeat the exercise to obtain 3 volatility portfolios. The summary statistics for these portfolios can be found in tables 2 and 3.

[**Place Table 2 about here**]

[**Place Table 3 about here**]

The cash flow betas are quite reasonable in size more stable across time. A strategy that goes long the high cash flow beta portfolio and short the low cash flow beta portfolio exhibits a return of approximately 3.5% in the full sample period and varies between 2.5% and 5% during the subsamples considered. Investors treat securities that have higher exposure to market real dividend growth as risky and demand a premium for holding them. On average, an increase of one unit in the cash flow growth beta is compensated by a return increase of 48 basis points.

Volatility carries a negative price of risk across all economic periods. A portfolio consisting of companies that do bad in times of high volatility earns on average returns that are 3.7%

higher compared to a portfolio of companies that act as a volatility hedge. The return differential is even higher at 4.85% during the last sample period which includes the recent financial crisis and the Great Recession that followed. Some care needs to be taken however in interpreting the subsample results for the volatility portfolios as it is harder to control for their inflation exposure and there is some overlap between the two type of risk premiums.

This caveat holds more generally. In the full sample the returns on the long-short portfolios with respect to cash flow and volatility risk are nearly orthogonal to each other, while their correlation with the returns on the inflation long-short portfolio returns 0.13 and -0.15 as can be seen in table 4. The correlations are slightly higher in absolute value across the three subperiods considered, reaching -0.34 between the cash flow and volatility strategies during the first subperiod and 0.28 and -0.24 for the correlation between these strategies and the inflation long-short portfolio during the last subsample.

[Place Table 4 about here]

The usual approach in the literature at this point would be to implement the second step of the Fama-MacBeth approach and use the time-series beta estimates in a cross sectional regression to pin down the prices for the three sources of risk. One potential concern with this approach relates to the accurate measurement of betas. In order to obtain a valid estimate of the price of risk, the ex-ante portfolio betas constructed from the time series regressions need to be good proxies for the "true" portfolio betas. It is well know for example in the context of market betas (see Black, Jensen, and Scholes (1972) and Fama and MacBeth (1973) as well as the references within) that, while the relative ranking of the ex-ante betas may be a good indication of the true beta relative ranking, the ex-ante beta range tends to overstate true beta range with low ex-ante portfolio betas underestimating the true ones and the high ex-ante portfolio betas overestimating the true values. More generally one might even express concerns with respect to factor selection and over fitting. In order to address these issues in the next section I write down a long run risk model with a built in inflation non-neutrality channel in the spirit of Piazzesi and Schneider (2007). The model, similar in structure to Bansal and Shaliastovich (2013), provides a motivation for the risk factors that I choose and has the advantage of pinning down the prices of risk for all the factors involved.

3 Theoretical Framework

The setup is that of a representative agent endowment economy with recursive preferences (Kreps and Porteus (1978), Epstein and Zin (1989), Weil (1989)) which allow for the separation of the relative risk aversion γ and intertemporal elasticity of substitution ψ parameters and, as a result, permit both of them to be higher than 1. Let

$$U_t = [(1-\delta)C_t^{1-\frac{1}{\psi}} + \delta(E_t U_{t+1}^{1-\gamma})^{\frac{1-\frac{1}{\psi}}{1-\gamma}}]^{\frac{1}{1-\frac{1}{\psi}}} \tag{3}$$

where U_t stands for the lifetime utility of the agent, C_t denotes time t consumption and δ is a subjective discount factor. For notational convenience let $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$. Whenever $\gamma > \frac{1}{\psi}$ the agent has a preference for early resolution of uncertainty. If $\gamma = \frac{1}{\psi}$ this utility function collapses to the standard CRRA case.

Epstein and Zin (1989) show that the log stochastic discount factor of the above economy can be written in terms of log consumption growth Δc_{t+1} and the return $r_{c,t+1}$ on the aggregate wealth portfolio that pays consumption as its dividend.

$$m_{t+1} = \theta log(\delta) - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta-1)r_{c,t+1} \tag{4}$$

Assuming the SDF and the return on the aggregate wealth are jointly log-normal the standard Euler equation

$$E_t[M_{t+1}R_{t+1}] = 1 \tag{5}$$

can be reframed as:

$$E_t\Delta c_{t+1} = \psi log\delta + \psi E_t r_{c,t+1} - \frac{\psi - 1}{\gamma - 1}V_t \tag{6}$$

where $V_t = \frac{1}{2}Var_t(m_{t+1} + r_{c,t+1})$ reflects the volatility of the SDF, of consumption growth and the covariance of the two.

Log-linearizing the budget constraint of this economy

$$W_{t+1} = (W_t - C_t)R_{C,t+1} \tag{7}$$

as it is standard in the literature produces

$$r_{c,t+1} = \kappa_0 + wc_{t+1} - \frac{1}{\kappa_1}wc_t + \Delta c_{t+1} \tag{8}$$

where $wc_t = log\frac{W_t}{C_t}$ is the log wealth to consumption ratio and κ_0 and κ_1 are the log-linearization parameters.

The model dynamics follow closely Piazzesi and Schneider (2007) and Bansal and Shaliastovich (2013). The consumption growth and inflation processes are specified exogenously allowing for persistent fluctuations in their conditional means and interactions of expected inflation and expected growth. Let $x_{c,t}$ and $x_{\pi,t}$ denote the expected consumption growth and expected inflation respectively. One has

$$\Delta c_{t+1} = \mu_c + x_{c,t} + \sigma_c \eta_{c,t+1} \tag{9}$$

$$\pi_{t+1} = \mu_\pi + x_{\pi,t} + \sigma_\pi \eta_{\pi,t+1} \tag{10}$$

where $\eta_{c,t+1}$ and $\eta_{\pi,t+1}$ are assumed to be i.i.d. N(0,1) and σ_c and σ_π are the conditional volatilities of the two processes.

The vector of conditional means $x_t = [x_{c,t}\ x_{\pi,t}]'$ follows a bivariate VAR(1) process with one source of time-varying conditional volatility σ_t.

$$x_{t+1} = \Pi x_t + \sigma_t \Sigma e_{t+1} \tag{11}$$

where $e_{t+1} = [e_{c,t+1}\ e_{\pi,t+1}]'$. I view $e_{c,t+1}$ and $e_{\pi,t+1}$ as structural shocks that each follow a N(0,1) distribution and are orthogonal to each other. I hence restrict the matrix

$$\Sigma = \begin{pmatrix} \varphi_c & 0 \\ 0 & \varphi_\pi \end{pmatrix} \tag{12}$$

to be diagonal. The interaction between the consumption growth and inflation in this economy is obtained by allowing the expected inflation and expected growth processes to directly affect each other. Consistent with this, the AR(1) matrix of loading is given by

$$\Pi = \begin{pmatrix} \rho_c & \rho_{c\pi} \\ \rho_{\pi c} & \rho_\pi \end{pmatrix} \tag{13}$$

Finally, for simplicity, I model the time varying conditional volatility process as an AR(1).

$$\sigma_{t+1}^2 = (1-\nu)\sigma^2 + \nu\sigma_t^2 + \sigma_w w_{t+1} \tag{14}$$

where w_{t+1} is distributed N(0,1), σ^2, ν and σ_w^2 are the mean, persistence and homoskedastic variance of the volatility process.

In the above economy one can show that the wealth to consumption ratio will be linear in the three state variables expected consumption growth, expected inflation and volatility.

$$wc_t = A_0 + A_{xc}x_{c,t} + A_{x\pi}x_{\pi,t} + A_\sigma\sigma_t^2 \tag{15}$$

where the loadings are given by:

$$A_{xc} = (1-\frac{1}{\psi})\chi_c,\ A_{x\pi} = (1-\frac{1}{\psi})\chi_{c\pi},\ A_\sigma = (1-\gamma)(1-\frac{1}{\psi})\chi_\sigma \tag{16}$$

and

$$\begin{aligned} \chi_c &= \frac{1-\kappa_1\rho_\pi}{(1-\kappa_1\rho_c)(1-\kappa_1\rho_\pi)-\kappa_1^2\rho_{\pi c}\rho_{c\pi}} \\ \chi_{c\pi} &= \frac{\kappa_1\rho_{c\pi}}{(1-\kappa_1\rho_c)(1-\kappa_1\rho_\pi)-\kappa_1^2\rho_{\pi c}\rho_{c\pi}} \\ \chi_\sigma &= \frac{1}{2(1-\kappa_1\nu)}(\kappa_1^2\chi_c^2\varphi_c^2 + \kappa_1^2\chi_{c\pi}^2\varphi_\pi^2) \end{aligned} \tag{17}$$

With an IES higher than 1 the substitution effect dominates the wealth effect and in response to positive news about the future growth the agent saves more driving up current prices and the wealth to consumption ratio. In the full sample data expected inflation has a negative effect on expected consumption growth (that is $\rho_{c\pi} < 0$)[6] implying a negative $\chi_{c\pi}$ estimate and $A_{x\pi} < 0$. If an increase in expected inflation forecasts lower future growth, then in response to a positive shock to expected inflation the agent will save less driving prices and the wealth to consumption ratio down. Finally if, in addition to the IES being higher than 1, the relative risk aversion coefficient is also higher than 1 a positive shock to volatility represents bad news for the economy and the wealth to consumption ratio will go down.

With the wealth to consumption ratio pinned down by the model it is trivial to solve for the return on wealth using the log-linearized budget constraint (8). The innovation to the real stochastic discount factor in this economy is then recovered from equation (4).

$$
\begin{aligned}
m_{t+1} - E_t m_{t+1} &= -\gamma \sigma_c \eta_{c,t+1} - (\gamma - \frac{1}{\psi})\kappa_1 \chi_c \varphi_c \sigma_t e_{c,t+1} \\
&- (\gamma - \frac{1}{\psi})\kappa_1 \chi_{c\pi} \varphi_\pi \sigma_t e_{\pi,t+1} - (1-\gamma)(\gamma - \frac{1}{\psi})\kappa_1 \chi_\sigma \sigma_w w_{t+1} \qquad (18)
\end{aligned}
$$

The price of risk for the current consumption shock is equal to the coefficient of relative risk aversion γ. With early resolution of uncertainty ($\gamma > \frac{1}{\psi}$), the price of expected consumption risk is positive. If inflation is non-neutral and has a negative effect on the real economy then the market price of expected inflation risk is negative. Finally, if in addition of early resolution of uncertainty you have $\gamma > 1$, then volatility also has a negative price of risk. This is consistent with low realized and expected consumption, high expected inflation and high volatility, all being the "bad" states of the world when the marginal utility of the representative agent is high.

The loadings $\kappa_1 \chi_c = \frac{\kappa_1(1-\kappa_1\rho_\pi)}{(1-\kappa_1\rho_c)(1-\kappa_1\rho_\pi)-\kappa_1^2\rho_{\pi c}\rho_{c\pi}}$ and $\kappa_1 \chi_{c\pi} = \frac{\kappa_1^2\rho_{c\pi}}{(1-\kappa_1\rho_c)(1-\kappa_1\rho_\pi)-\kappa_1^2\rho_{\pi c}\rho_{c\pi}}$ are noth-

[6] Also typically $\rho_{\pi c} > 0$.

ing but the present value effect of contemporaneous persistent consumption and inflation shocks on future path of expected consumption growth. Indeed, the stochastic discount factor above can be easily recast in the language of Bansal et al. (2014) and Campbell et al. (2015). The news about current and future consumption growth are given by

$$\widetilde{N_{C,t+1}} = \Delta c_{t+1} - E_t(\Delta c_{t+1}) = \sigma_c \eta_{c,t+1} \tag{19}$$

and

$$\begin{aligned} N_{ECF,t+1} = \sum_{j=1}^{\infty} \kappa_1^j (E_{t+1} - E_t) \Delta c_{t+1+j} &= \frac{\kappa_1 (1 - \kappa_1 \rho_\pi)}{(1 - \kappa_1 \rho_c)(1 - \kappa_1 \rho_\pi) - \kappa_1^2 \rho_{c\pi} \rho_{\pi c}} \varphi_c \sigma_t e_{c,t+1} + \\ &+ \frac{\kappa_1^2 \rho_{c\pi}}{(1 - \kappa_1 \rho_c)(1 - \kappa_1 \rho_\pi) - \kappa_1^2 \rho_{c\pi} \rho_{\pi c}} \varphi_\pi \sigma_t e_{\pi,t+1} \\ &= \kappa_1 \chi_c \varphi_c \sigma_t e_{c,t+1} + \kappa_1 \chi_{c\pi} \varphi_\pi \sigma_t e_{\pi,t+1} \end{aligned} \tag{20}$$

As it is obvious from above, if consumption growth and inflation affect each other in equilibrium, the typical news term about future consumption will depend on both the consumption and the inflation long-run shocks. Under the null of the model this term can then be broken down into two separate components that have prices of risk with opposite signs and that carry very different implications for the pricing of assets. From this point on, I will refer to the present value impact of the expected inflation shock as "news about expected inflation" and I will call the residual term (which is nothing but the present value impact of the expected consumption shock) "news about expected cash flows", that is:

$$\widetilde{N_{ECF,t+1}} = \chi_c \varphi_c \sigma_t e_{c,t+1}, \quad \widetilde{N_{E\pi,t+1}} = |\chi_{c\pi}| \varphi_\pi \sigma_t e_{\pi,t+1}, \quad N_{ECF} = \widetilde{N_{ECF,t+1}} - \text{sign}(\chi_{c\pi}) \widetilde{N_{E\pi,t+1}} \tag{21}$$

An observation is in order. By taking the absolute value of $\chi_{c\pi}$ I ensure that the news about expected inflation term has the same direction as the expected inflation shock $e_{\pi,t+1}$. The sign of $\chi_{c\pi}$ is pulled into the price of risk which becomes $(\gamma - \frac{1}{\psi})\,\mathrm{sign}(\chi_{c\pi})$. In the full sample, since inflation is bad news for future growth ($\rho_{c\pi} < 0$), this leads to $\chi_{c\pi} < 0$ and a negative price of risk for expected inflation. If inflation becomes pro-cyclical, as might be the case starting in the late 90's, then the signs of $\rho_{c\pi}$ and $\chi_{c\pi}$ would flip leading to a positive price of expected inflation risk.

It is important to note that under the null of the model the two components of $N_{ECF,t+1}$ are orthogonal to each other so that the present value effect of expected inflation shocks ($\chi_{c\pi}$) can be read from the following restriction

$$\chi_{c\pi} = \frac{Cov(N_{ECF,t+1}, \varphi_\pi \sigma_t e_{\pi,t+1})}{Var(\varphi_\pi \sigma_t e_{\pi,t+1})} \tag{22}$$

It turns out that, in the model, the volatility term in the log Euler equation is linear in the conditional volatility of the wealth return

$$V_t = \frac{1}{2} Var_t(m_{t+1} + r_{c,t+1}) = const. + \frac{\theta^2}{2} Var_t(N_{R,t+1}) \tag{23}$$

so that news about volatility are a scaled version of news about the wealth return conditional variance with the proportionality constant being equal to $\frac{\theta^2}{2}$

$$\widetilde{N_{V,t+1}} = \frac{\theta^2}{2} N_{Var_t(N_{R,t+1})} = (1 - \frac{1}{\psi})^2 \kappa_1 \chi_\sigma \sigma_w w_{t+1} \tag{24}$$

where news about volatility are defined as $\widetilde{N_{V,t+1}} = \sum\limits_{j=1}^{\infty} \kappa_1^j (E_{t+1} - E_t) V_{t+j}$.

Putting everything together I can express the shock to the real stochastic discount factor in terms of news about contemporaneous and expected consumption growth, news about expected inflation and news about volatility.

$$N_{m,t+1} = -\lambda_C \widetilde{N_{C,t+1}} - \lambda_{ECF} \widetilde{N_{ECF,t+1}} - \lambda_{E\pi} \widetilde{N_{E\pi,t+1}} - \lambda_V \widetilde{N_{V,t+1}} \tag{25}$$

with

$$\lambda_C = \gamma,\ \lambda_V = -\frac{(\gamma-1)(\gamma-\frac{1}{\psi})}{(1-\frac{1}{\psi})^2}$$
$$\lambda_{ECF} = (\gamma - \frac{1}{\psi}),\ \lambda_{E\pi} = \text{sign}(\chi_{c\pi})(\gamma - \frac{1}{\psi}) \tag{26}$$

The market price of news about contemporaneous consumption is still γ. News about expected cash flows and news about expected inflation carry prices of risk equal to $\gamma - \frac{1}{\psi}$ and $(\gamma - \frac{1}{\psi})\,\text{sign}(\chi_{c\pi})$ respectively. As mentioned before, if the representative agent has a preference for early resolution of uncertainty then price of risk for news about expected cash flows risk is positive and the price of risk for expected inflation news is negative if expected inflation has a negative impact on the future consumption path. If in addition to early resolution of uncertainty $\gamma > 1$ then news about volatility carry a negative price of risk equal to $-\frac{(\gamma-1)(\gamma-\frac{1}{\psi})}{(1-\frac{1}{\psi})^2}$.

Using the log Euler equations (6) for a generic asset with return $r_{i,t+1}$ and for the risk free rate $r_{f,t}$ one can express the risk premium on asset i as the negative of the covariance of the log SDF with $r_{i,t+1}$

$$E_t r_{i,t+1} - r_{f,t} + \frac{1}{2} Var_t(r_{i,t+1}) = -Cov_t(m_{t+1}, r_{i,t+1}) \tag{27}$$

which can further be decomposed into risk compensation for each of the news terms.

$$
\begin{aligned}
E_t r_{i,t+1} - r_{f,t} + \frac{1}{2} Var_t(r_{i,t+1}) &= \lambda_C Cov_t(\widetilde{N_{C,t+1}}, r_{i,t+1}) + \lambda_{ECF} Cov_t(\widetilde{N_{ECF,t+1}}, r_{i,t+1}) \\
&+ \lambda_{E\pi} Cov_t(\widetilde{N_{E\pi,t+1}}, r_{i,t+1}) + \lambda_V Cov_t(\widetilde{N_{V,t+1}}, r_{i,t+1}) \quad (28)
\end{aligned}
$$

4 Empirical Results

The wealth to consumption ratio wc_t and the return on wealth $r_{c,t+1}$ are not directly observable in the data. As a result I will follow the large body of literature (starting with Epstein and Zin (1991)) and assume that the wealth portfolio is well proxied by a broad stock market index. As a result I will interpret the wealth to consumption ratio as the price dividend ratio, the return on wealth as the return on the market and consumption growth and the market dividend growth adjusted for inflation.

The most parsimonious VAR one can write down must contain the real market dividend growth Δd_t, log market price-dividend ratio pd_t, expected inflation $E_t[\pi_{t+1}]$, and realized market volatility RV_t. I add the ex-post real risk free rate $rf_{r,t}$ as the model also has implications for this variable and adding it can improve predictability[7]. The real dividend growth and the price dividend ratio are constructed from the CRSP returns on the NYSE, NASDAQ and AMEX portfolio with and without dividends. One year ahead expected inflation comes from the Philadelphia Fed Survey of Professional Forecasters. The corresponding inflation data comes from BEA tables. Realized volatility for each year is constructed as the sum of the 4 squared quarterly inflation adjusted market returns for that year. A measure of the ex-post real risk free rate is constructed as the difference between the yield on the 1-year

[7]Results are robust if the ex-post real risk free rate is left out.

Fama-Bliss bond and my inflation measure. All growth rates and returns are continuously compounded and logs have been taken where appropriate. Frequency is annual and the sample period runs from 1970 to 2014. As opposed to section 2, I lose one year because the one year ahead expected inflation forecast is not available for 1969.

Denoting by z_t the vector assumed to capture the state of the economy and restricting the dynamics to a VAR(1) we have

$$z_t = [\Delta d_t \ E_t[\pi_{t+1}] \ RV_t \ pd_t \ rf_{r,t}]' \tag{29}$$

where

$$z_{t+1} = \phi_0 + \Phi z_t + u_{t+1}$$

and $u_{t+1} \sim N(0, \Sigma_u)$.

Further let $\iota_d = [1\ 0\ 0\ 0\ 0]'$, $\iota_\pi = [0\ 1\ 0\ 0\ 0]'$ and $\iota_v = [0\ 0\ 1\ 0\ 0]'$ be selection vectors and $Q = \kappa_1 \Phi (I - \kappa_1 \Phi)^{-1}$ be the long-run response matrix for this VAR.

The classic news terms $N_{C,t+1}$, $N_{ECF,t+1}$ can be read from the VAR in the usual way

$$\widetilde{N_{C,t+1}} = \iota_d' u_{t+1} \tag{30}$$

$$N_{ECF,t+1} = \iota_d' Q u_{t+1} \tag{31}$$

and so can the shock to expected inflation

$$\varphi_\pi \sigma_t e_{\pi,t+1} = \iota_\pi' u_{t+1} \tag{32}$$

Combining the two and using the theoretical restriction (22) identifies

$$\chi_{c\pi} = \frac{\iota_d' Q \Sigma_u \iota_\pi}{\iota_\pi' \Sigma_u \iota_\pi} \tag{33}$$

The news about expected cash flows and news about expected inflation are now trivial to construct

$$\widetilde{N_{E\pi,t+1}} = \left| \frac{\iota_d' Q \Sigma_u \iota_\pi}{\iota_\pi' \Sigma_u \iota_\pi} \right| \iota_\pi' u_{t+1} \tag{34}$$

$$\widetilde{N_{ECF,t+1}} = \iota_d' Q u_{t+1} - \text{sign}\left(\frac{\iota_d' Q \Sigma_u \iota_\pi}{\iota_\pi' \Sigma_u \iota_\pi} \right) \left| \frac{\iota_d' Q \Sigma_u \iota_\pi}{\iota_\pi' \Sigma_u \iota_\pi} \right| \iota_\pi' u_{t+1} \tag{35}$$

Finally, I proxy the conditional volatility of the return on wealth $Var_t(r_{c,t+1})$ by the realized variance of the market return RV_t. This allows me to read the news component $N_{Var_t(r_{c,t+1})}$ in (24) directly from the VAR implying that

$$N_{V,t+1} = \frac{\theta^2}{2} \kappa_1 \iota_v' (I + Q) u_{t+1} \tag{36}$$

Equations (30), and (34) - (36) completely pin down the model implied risk factors for this economy.

The VAR is estimated by first demeaning the variables and then running equation by equation OLS. Results can be found in table 5. Newey-West standard errors with 4 lags are shown in parentheses below each estimate. Given the "kitchen sink" nature of the VAR it is hard to comment on individual coefficients. It is worth emphasizing however that, in the given sample, expected inflation is quite predictable with an R^2 of almost 87% and displays rather strong dependence on its lagged value with a coefficient of 0.89 that is highly statistically significant. The real risk free rate also helps forecast expected inflation in the data. Furthermore expected inflation predicts future dividend growth and the future price dividend ratio with a negative coefficient that is also statistically significant in the second case.

[Place Table 5 about here]

The full sample VAR estimates imply a negative value of $\chi_{c\pi}$ equal to -8.34. A shock to expected inflation constitutes bad news for economy. One percentage point increase in expected inflation leads to a drop in future expected consumption growth of 8.34% in present value. The standard deviation of expected inflation shocks over the period 1970 to 2014 is 0.72%. A one standard deviation shock to expected inflation would therefore lead to a drop in expected future consumption growth of approximately 6% in present value.

If expected inflation shocks have such a high impact on expected cash flows, it stands to reason that positive expected inflation news should be accompanied by an increase in the risk premium. In long run risk models the time variation in risk premia is generated by time varying volatility (indeed Bansal et al. (2014) show that discount rate news and volatility news are strongly correlated with one another) so would would expect to see a positive correlation between expected inflation and volatility news. This intuition is confirmed by the data as the correlation between the two news factors is approximately 42% during the sample period considered. One should note the result above recovers, in my setup, the finding of Bekaert and Engstrom (2010) that expected inflation and risk premium have a correlation that plateaus at approximately 40% in the long run.

I test the cross sectional implications of the model by pricing the excess returns on a wide range of assets consisting of the market portfolio, ten industry portfolios taken from Kenneth French's website, the Fama-Bliss discount bonds as well as a spot commodity index constructed by the Commodity Research Bureau. The ten Fama-French portfolios are broadly representative of industries producing durable and nondurable goods, manufacturing, energy, high tech, telecommunications, retail, healthcare and utilities as well as a catch all portfolio for all remaining firms that were not included in any of the previous categories. I chose to focus on industry portfolios rather than portfolios sorted on other dimensions (like size or book to market), because I expect there to be more variation in terms of inflation exposure across industries which will makes them a more interesting test asset from the perspective of this paper. For commodities I use the CRB BLS commodity spot index because of its long time span.[8] The index is a measure of the spot price changes in 22 commodities whose markets are assumed to be the first to respond to a change in economic conditions. The set of commodities used is composed of raw materials[9] making up about 60% of the index and foodstuffs[10] accounting for the remaining 40%. The index is typically viewed as a general economic indicator used for gauging the direction of prices. An increase in price for the basic commodities underlying the index could translate into higher production costs throughout various sectors of the economy and higher inflation. An economic agent that holds the basket of commodities comprising the index on the other hand would see a larger return on his portfolio and thus be hedged from the increase in inflation.

For each of these portfolios I first extract the shocks by running a time series regression of the portfolio real return on the vector of predictive variables z_t. With the news components pinned down by the VAR estimates, I then compute the time series covariances between the portfolio shocks and the the model implied factors. The prices of risk (which depend only

[8]The index is constructed as the geometric mean of price ratios of the commodity prices today to the commodity prices on the base date, multiplied by 100. The log change in the index will therefore reflect an equal weighted average of the log changes in each commodity's price.

[9]Including burlap, copper scrap, cotton, hides, lead scrap, print cloth, rosin, rubber, steel scrap, tallow, tin, wool tops, and zinc.

[10]Including butter, cocoa beans, corn, cottonseed oil, hogs, lard, steers, sugar, and wheat.

on the preference function parameters γ and ψ) can then in principle be pinned down in a cross sectional regression that has the expected returns of the 16 considered portfolios on the left hand side and their covariances with the news components as explanatory variables. The first order conditions for gamma and psi given by a nonlinear least squares estimator lead to a system of nonlinear equations that can be hard to solve. As a consequence I take a simple two step approach to pin down the relative risk aversion and IES parameters. First, for a given value of ψ, the first order condition with respect to γ reduces to a third order equation that can be easily solved. The solution that minimizes the sum of squared errors is always higher than 1 and higher than $\frac{1}{\psi}$ ensuring early resolution of uncertainty, a positive price for expected cash flow risk and a negative price for the expected inflation and volatility risks. Second, I let the IES parameter evolve between 0.5 and 3 with a step of 0.05. For each value of ψ I perform the procedure described in the first step. I then pick the IES that achieves the lowest sum of squared errors and the corresponding relative risk aversion value. The coefficients identified by this procedure are $\psi = 2.2$ and $\gamma = 4.13$. The fit generated by these parameter values can be seen in tables 6A and 6B.

[Place Table 6A about here]

[Place Table 6B about here]

The model matches the overall premiums well. The inflation premium on the market portfolio is a 1.5% compared to approximately 2% and 3% for cash flow growth and volatility risk emphasizing once more the importance of inflation risk for asset markets and the economy. There is also strong variation in the inflation premium across the industry portfolios. The retail sector has the strongest exposure to inflation risk with a risk premium of approximately 2.7% and is followed closely by the durable goods sector where the premium is 2.55%. The inflation premium for nondurable goods is lower at 1.95% consistent with the findings of Eraker, Shaliastovich, and Wang (2015) who document that the negative impact of expected inflation on future consumption growth is even stronger in the case of durable goods consumption as opposed to nondurables. At the opposite end one finds the energy portfolio that has a negative inflation premium of -0.33%. This finding is quite reasonable if one considers that the energy portfolio is to a large extend driven by oil companies. If the price of oil goes up, this generates a surge in prices through the entire economy and a rise in inflation rates. While the increase in inflation is bad news for most economic agents, this is precisely the state of the world where oil companies are doing well. The energy portfolio can therefore be seen as an inflation hedge which is reflected in the premium I find.

The model identifies an inflation premium in nominal bonds (consistent with a vast literature on this topic) that is increasing in the maturity of the bond. Indeed, a nominal bond is offering a fixed money payment at some date in the future. The further away that date is, the lower its value in real terms if inflation is expected to be high. Moreover, if periods of high expected inflation correspond in the data with states of low expected economic growth (to which a long-run risk agent would be very averse) a payment that is already lower in real term comes at the worse possible moment. Hence the longer the maturity, the higher the inflation premium. Finally, the commodity index displays a negative inflation premium of -1.11% consistent with the commonly circulated idea that commodities act a hedge against inflation. Indeed, the commodities that make up the index are supposed to be foods and materials that are at the base of the production chain. An increase in price for these products will drive up the cost of other fabricated goods and reverberate through the economy. Holding these commodities hedges the owner from the adverse effect of increased inflation in a manner similar to that discussed above in the case of oil.

One might wonder why the focus on industry portfolios as opposed to the size and book to market sorted portfolios as it is typical in the literature. Table 6C shows the model fit for the five size quintile portfolios and the five quintile book to market sorted portfolios (taken once more from Kenneth French's website and aggregated to annual frequency). For consistency reasons I keep the same values for the relative risk aversion and intertemporal elasticity substitution coefficients as before. As can be seen the model provides a good fit here as well and generates size and book to market spreads of the correct sign and similar values to those found in the data. However there is no clear pattern in terms of inflation risk premia.

[**Place Table 6C about here**]

Finally, it would be interesting to look at how the price of expected inflation risk evolves over time. Ideally one would have a rolling window and within each window reestimate/recalibrate the entire model. Due to the low frequency of the data and the small overall number of observations this approach is unfeasible. The VAR estimates for example would be poorly estimated and quite unstable even for a window of 10-15 years. Moreover the risk premium estimates for the portfolios I match would also be imprecisely estimated over such a short span. As a compromise I will keep the full sample VAR estimates as well as the shocks to the state variable vector and matched portfolios based on these estimates. Furthermore I stick to preference parameters previously chosen (i.e. $\gamma = 4.13$ and $\psi = 2.2$). For these values I use a 5-year rolling window to recompute the $\chi_{c\pi}$ parameter as well as the variance of the news components over each such interval. The expected inflation price of risk generated by this approach is plotted in figure 3.

$$\widetilde{\lambda_{E\pi,t}} = \text{sign}(\chi_{c\pi,t})(\gamma - \frac{1}{\psi})Var_t(\widetilde{N_{E\pi,t+1}}) \tag{37}$$

[Place Figure 3 about here]

The price of expected inflation risk is large and strongly negative both in the mid 70's and in the early 80's, a timing consistent with the two oil shocks that hit the US economy and the recessions they have generated. During the 90's, a period where inflation has been tamed, the price of expected inflation risk is almost constant but still negative. With long run inflation expectations well anchored at around 2% and little volatility in inflation expectations in general, the price of inflation risk gets close to zero during the last the decade and it briefly crosses into positive territory during the the Great Recession. This is consistent with the results from the portfolio sorts that find a much smaller and slightly negative spread for the long-short portfolio giving exposure to expected inflation risk.

5 Conclusion

This paper shows that expected inflation risk is priced in the cross section of US stock returns while controlling for cash flow growth and volatility risk. This finding is interpreted through the prism of an inflation non-neutrality channel that allows expected inflation to directly affect the marginal utility of the representative agent. Embedding such a channel in the standard long run risk model leads to a decomposition of the real stochastic discount factor into news about short and long run cash flow growth, news about expected inflation and news about volatility. The model does a good job at simultaneously matching the risk premia on a wide cross section of assets including industry portfolios, nominal bonds and a commodity index. The cross sectional and time series variation in the inflation risk premium is then investigated. The model identifies bonds as being exposed to inflation risk and commodities as a a hedge. The energy industry which is largely formed of oil companies has a small and negative expected inflation risk premium while the same premia is large for industries like retail. Consistent with the idea that high expected inflation has a larger impact on future durable goods consumption as opposed to nondurables, the model identifies a larger inflation premium in the former industry portfolio. The time-varying price of risk for expected inflation is large in 70's and 80's when inflation was rampant and has come down a lot since. Over the past decade the inflation price of risk has been close to zero and has occasionally flipped sign crossing into positive territory especially during the recent Great Recession.

All the results above come from a market based approach where the return on wealth in the model is identified in the data as a stock market return. A natural and interesting extension would be to bring labor income growth to the table and model it jointly with consumption growth and inflation. This would allow me to study the interaction between expected inflation and the return on human capital as well as make broader statements about the impact of inflation on the wealth to consumption ratio and true return on wealth. Finally such an approach could naturally lead to estimates of the "cost" of inflation risk, specifically how much consumption would an agent be willing to sacrifice in order to avoid an increase in inflation.

References

Martin J. Bailey. The welfare cost of inflationary finance. Journal of Political Economy, 64: 93–110, 1956.

Ravi Bansal and Ivan Shaliastovich. A long-run risks explanation of predictability puzzles in bond and currency markets. Review of Financial Studies, 26:1–33, 2013.

Ravi Bansal, Dana Kiku, Ivan Shaliastovich, and Amir Yaron. Volatility, the macroeconomy, and asset prices. The Journal of Finance, 69:2471–2511, 2014.

Geert Bekaert and Eric Engstrom. Inflation and the stock market: Understanding the fed model. Journal of Monetary Economics, 57:278–294, 2010.

Geert Bekaert and Xiaozheng Wang. Inflation risk and the inflation risk premium. Economic Policy, 25:755–806, 2010.

Fisher Black, Michael C. Jensen, and Myron Scholes. The capital asset pricing model: Some empirical tests. Studies in the Theory of Capital Markets, pages 79–121, 1972.

Zvi Bodie. Common stocks as a hedge against inflation. The Journal of Finance, 31:459–470, 1976.

John Y. Campbell, Adi Sunderam, and Luis M. Viceira. Inflation bets or deflation hedges? the changing risks of nominal bonds. Working Paper, 2013.

John Y. Campbell, Stefano Giglio, Christopher Polk, and Robert Turley. An intertemporal CAPM with stochastic volatility. Working Paper, 2015.

Thomas F. Cooley and Gary D. Hansen. The inflation tax in a real business cycle model. American Economic Review, 79:733–748, 1989.

Michael Dotsey and Peter Ireland. The welfare cost of inflation in general equilibrium. Journal of Monetary Economics, 37:29–47, 1996.

Fernando M. Duarte. Inflation risk and the cross section of stock returns. Working Paper, 2013.

Larry G. Epstein and Stanley Zin. Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica, 57:937–969, 1989.

Larry G. Epstein and Stanley Zin. Substitution, risk aversion and the temporal behavior of consumption and asset returns: An empirical analysis. Journal of Political Economy, 99: 263–286, 1991.

Bjorn Eraker, Ivan Shaliastovich, and Wenyu Wang. Durable goods, inflation risk, and equilibrium asset prices. Working Paper, 2015.

Eugene F. Fama and Kenneth R. French. The cross-section of expected stock returns. The Journal of Finance, 47:427–465, 1992.

Eugene F. Fama and James D. MacBeth. Risk, return and equilibrium: Empirical tests. Journal of Political Economy, 81:607–636, 1973.

Eugene F. Fama and G. William Schwert. Asset returns and inflation. The Journal of Financial Economics, 5:115–146, 1977.

Stanley Fisher. Towards an understanding of the costs of inflation: Ii. Carnegie-Rochester Conference Series on Public Policy 15, pages 5–41, 1981.

Milton Friedman. The optimum quantity of money. The Optimum Quantity of Money, and Other Essays, 1969.

David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. Econometrica, 46:185–200, 1978.

Howard Kung. Macroeconomic linkages between monetary policy and the term structure of interest rates. Journal of Financial Economics, 115:42–57, 2015.

Ricardo Lagos and Randall Wright. A unified framework for monetary theory and policy analysis. Journal of Political Economy, 113:463–484, 2005.

Robert E. Jr. Lucas. Discussion of: Stanley fischer, 'towards an understanding of the costs of inflation: Ii'. Carnegie-Rochester Conference Series on Public Policy 15, pages 43–52, 1981.

Robert E. Jr. Lucas. Inflation and welfare. Econometrica, 68:247–274, 2000.

Monika Piazzesi and Martin Schneider. Equilibrium yield curves. NBER Macroeconomics Annual 2006, pages 389–442, 2007.

William G. Schwert. The adjustment of stock prices to information about inflation. The Journal of Finance, 36:15–29, 1981.

Michael Weber. Nominal rigidities and asset pricing. Working Paper, 2015.

Philippe Weil. The equity premium puzzle and the risk free rate puzzle. Journal of Monetary Economics, 24:401–421, 1989.

6 Tables

Table 1: Portfolios Exposed to Expected Inflation Risk

	β_π	β_{CF}	β_V	$E[R^e]$	% MV
		1974 Q2 - 2014 Q4			
Low β_π	-56.88	1.23	-4.32	10.88%	26.46%
Med β_π	-12.47	1.24	-4.37	9.83%	44.09%
High β_π	28.44	1.28	-4.71	8.84%	29.45%
(Low β_π)-(High β_π)				2.04%	
		1974 Q2 - 1987 Q4			
Low β_π	-13.46	1.73	-4.52	11.19%	27.40%
Med β_π	5.07	1.77	-5.33	8.90%	43.96%
High β_π	23.58	1.87	-6.46	5.77%	28.64%
(Low β_π)-(High β_π)				5.42%	
		1988 Q1 - 1997 Q4			
Low β_π	-55.42	1.35	-0.72	15.35%	21.48%
Med β_π	-17.14	1.23	-1.42	13.73%	44.65%
High β_π	16.81	1.09	-2.25	13.04%	33.87%
(Low β_π)-(High β_π)				2.31%	
		1998 Q1 - 2014 Q4			
Low β_π	-92.86	0.77	-6.28	8.12%	28.62%
Med β_π	-23.92	0.81	-5.32	8.35%	43.87%
High β_π	39.22	0.92	-4.73	8.85%	27.50%
(Low β_π)-(High β_π)				-0.73%	

The table presents summary statistics over various sample periods for portfolios that differ in their exposure to expected inflation risk (Low β_π, Med β_π, and High β_π), but have roughly similar exposures to cash flow growth and volatility risk. For each period the last line represents a long-short strategy where you buy the portfolio that has low (negative) inflation exposure and sell the portfolio that has high (positive) exposure. The first three columns are the average ex-ante betas of these portfolios with respect to expected inflation, cash flow growth and volatility risks. The fourth column gives the average annualized excess return on each portfolio for the quarter following the portfolio formation period. The last column is the average market value of the portfolio as a percentage of total market value.

Table 2: Portfolios Exposed to Cash Flow Growth Risk

	β_π	β_{CF}	β_V	$E[R^e]$	% MV
1974 Q2 - 2014 Q4					
Low β_{CF}	-13.18	-2.40	-5.94	8.26%	33.72%
Med β_{CF}	-11.07	1.12	-4.02	9.49%	44.32%
High β_{CF}	-16.65	5.03	-3.43	11.82%	21.96%
(High β_{CF})-(Low β_{CF})				3.56%	
1974 Q2 - 1987 Q4					
Low β_{CF}	5.51	-3.26	-5.86	6.62%	36.70%
Med β_{CF}	4.63	1.60	-5.17	8.58%	42.72%
High β_{CF}	5.04	7.04	-5.29	10.64%	20.58%
(High β_{CF})-(Low β_{CF})				4.02%	
1988 Q1 - 1997 Q4					
Low β_{CF}	-8.85	-2.24	-5.35	11.10%	30.33%
Med β_{CF}	-15.76	1.05	-1.17	14.98%	45.83%
High β_{CF}	-31.14	4.87	2.13	16.06%	23.84%
(High β_{CF})-(Low β_{CF})				4.96%	
1998 Q1 - 2014 Q4					
Low β_{CF}	-30.85	-1.79	-6.35	7.91%	33.30%
Med β_{CF}	-21.02	0.79	-4.77	7.10%	44.72%
High β_{CF}	-25.68	3.49	-5.21	10.33%	21.98%
(High β_{CF})-(Low β_{CF})				2.42%	

The table presents summary statistics over various sample periods for portfolios that differ in their exposure to cash flow growth risk (Low β_{CF}, Med β_{CF}, and High β_{CF}), but have roughly similar exposures to cash flow growth and volatility risk. For each period the last line represents a long-short strategy where you buy the portfolio that has high (positive) cash flow exposure and sell the portfolio that has low (negative) exposure. The first three columns are the average ex-ante betas of these portfolios with respect to expected inflation, cash flow growth and volatility risks. The fourth column gives the average annualized excess return on each portfolio for the quarter following the portfolio formation period. The last column is the average market value of the portfolio as a percentage of total market value.

Table 3: Portfolios Exposed to Volatility Risk

	β_π	β_{CF}	β_V	$E[R^e]$	% MV
		1974 Q2 - 2014 Q4			
Low β_V	-20.89	1.14	-14.01	11.84%	21.94%
Med β_V	-11.66	1.24	-4.41	9.56%	43.37%
High β_V	-8.36	1.37	5.02	8.16%	34.69%
(Low β_V)-(High β_V)				3.68%	
		1974 Q2 - 1987 Q4			
Low β_V	16.15	1.77	-14.13	10.18%	21.47%
Med β_V	4.20	1.78	-5.42	8.27%	38.68%
High β_V	-5.16	1.83	3.24	7.36%	39.85%
(Low β_V)-(High β_V)				2.82%	
		1988 Q1 - 1997 Q4			
Low β_V	-16.51	0.91	-11.10	15.67%	23.40%
Med β_V	-14.85	1.18	-1.64	13.56%	48.08%
High β_V	-24.37	1.59	8.35	12.89%	28.52%
(Low β_V)-(High β_V)				2.78%	
		1998 Q1 - 2014 Q4			
Low β_V	-53.42	0.76	-15.62	10.96%	21.47%
Med β_V	-22.60	0.85	-5.21	8.29%	44.39%
High β_V	-1.53	0.88	4.50	6.11%	34.14%
(Low β_V)-(High β_V)				4.85%	

The table presents summary statistics over various sample periods for portfolios that differ in their exposure to volatility risk (Low β_V, Med β_V, and High β_V), but have roughly similar exposures to cash flow growth and expected inflation risk. For each period the last line represents a long-short strategy where you buy the portfolio that has low (negative) volatility exposure and sell the portfolio that has high (positive) exposure. The first three columns are the average ex-ante betas of these portfolios with respect to expected inflation, cash flow growth and volatility risks. The fourth column gives the average annualized excess return on each portfolio for the quarter following the portfolio formation period. The last column is the average market value of the portfolio as a percentage of total market value.

Table 4: Correlation Structure for the Long-Short Portfolios

1974 Q2 - 2014 Q4	Infl	CF	Vol
Infl	1.00	0.13	-0.15
CF	0.13	1.00	-0.05
Vol	-0.15	-0.05	1.00

1974 Q2 - 1987 Q4	Infl	CF	Vol
Infl	1.00	-0.08	-0.10
CF	-0.08	1.00	-0.34
Vol	-0.10	-0.34	1.00

1988 Q1 - 1997 Q4	Infl	CF	Vol
Infl	1.00	0.18	0.11
CF	0.18	1.00	0.10
Vol	0.11	0.10	1.00

1998 Q1 - 2014 Q4	Infl	CF	Vol
Infl	1.00	0.28	-0.24
CF	0.28	1.00	0.21
Vol	-0.24	0.21	1.00

The table presents the correlations during various sample periods between the returns on the long-short portfolios that give exposure to expected inflation, cash flow growth and volatility risk respectively. For expected inflation and volatility the long-short strategy is constructed by going long the portfolio having a low (negative) exposure to the risk source and short the portfolio having a high (positive) exposure. For cash flow risk the opposite is true, one buys the high cash flow beta portfolio and shorts the low cash flow beta one.

Table 5: Market VAR Estimates

	Δd_t	$E_t[\pi_{t+1}]$	RV_t	pd_t	$rf_{r,t}$	R^2
Δd_{t+1}	0.17	-0.39	-0.11	0.01	-0.66	13.26%
	(0.14)	(1.01)	(0.45)	(0.05)	(0.46)	
$E_{t+1}[\pi_{t+2}]$	0.00	0.89	-0.04	0.00	-0.08	86.79%
	(0.01)	(0.10)	(0.05)	(0.00)	(0.03)	
RV_{t+1}	-0.02	0.32	0.21	0.02	0.17	8.62%
	(0.04)	(0.38)	(0.13)	(0.01)	(0.13)	
pd_{t+1}	0.03	-4.41	1.22	0.75	0.84	84.01%
	(0.28)	(1.83)	(0.99)	(0.11)	(0.90)	
$rf_{r,t+1}$	0.01	0.36	-0.17	0.01	0.81	76.47%
	(0.02)	(0.10)	(0.06)	(0.00)	(0.04)	

Annual frequency multivariate VAR(1) for the real dividend growth, price to dividend ratio, expected inflation, realized volatility, and real risk free rate. I first demean the variables then estimate the VAR using equation by equation OLS. Newey West standard errors with 4 lags are shown in parenthesis below each estimate. The real dividend growth and the price dividend ratio are constructed from the CRSP returns on the NYSE, NASDAQ and AMEX portfolio with and without dividends. One year ahead expected inflation comes from the Philadelphia Fed Survey of Professional Forecasters. I use the change in the GDP price index as my measure of inflation because the SPF data for this measure goes back further. The corresponding inflation data (used in obtaining real versions of the variables) comes from BEA tables. Realized volatility for each year is constructed as the sum of the 4 quarterly squared real market returns for that year. A measure of the ex-post real risk free rate is constructed as the difference between the yield on the 1-year Fama-Bliss bond and my inflation measure. All growth rates and returns are continuously compounded and logs have been taken where appropriate. Frequency is annual and the sample period runs from 1970 to 2014.

Table 6A: Market and Industry Portfolios Risk Premia

	Mkt	NonDur	Dur	Manuf	Energy	HTech
Data	5.97%	8.11%	6.78%	6.61%	7.68%	6.39%
Model	6.49%	6.39%	8.65%	5.92%	3.59%	6.38%
Cash Flow	1.96%	0.96%	2.13%	1.51%	1.55%	2.44%
Inflation	1.48%	1.95%	2.55%	1.28%	-0.33%	1.93%
Volatility	3.06%	3.48%	3.98%	3.12%	2.37%	2.02%
		Telcm	Shops	Health	Utils	Other
Data		7.14%	7.42%	7.37%	6.22%	6.14%
Model		7.04%	7.59%	3.72%	6.67%	8.75%
Cash Flow		2.81%	1.38%	0.68%	1.02%	1.92%
Inflation		1.94%	2.68%	1.45%	1.61%	1.73%
Volatility		2.30%	3.53%	1.60%	4.04%	5.10%
Model		λ_C	λ_{ECF}	$\lambda_{E\pi}$	λ_V	
Implied		γ	$\gamma - \frac{1}{\psi}$	$-(\gamma - \frac{1}{\psi})$	$-\frac{(\gamma-1)(\gamma-\frac{1}{\psi})}{(1-\frac{1}{\psi})^2}$	
Risk Price		4.13	3.67	-3.67	-38.61	

The table shows the model fit for the market portfolio as well as 10 industry portfolios taken from Kenneth French's website. The first row is the excess return on the portfolio in the data. The second row shows the model implied risk premium for each portfolio. The following three rows give a breakdown of the model implied risk premium into 3 components (cash-flow risk, inflation risk and volatility risk) where cash-flow risk premium is defined as $\lambda_C \beta_{i,C} + \lambda_{ECF} \beta_{i,ECF}$, expected inflation risk premium as $\lambda_{E\pi} \beta_{i,E\pi}$ and volatility premium as $\lambda_V \beta_{i,V}$. The model implied market prices of risk are displayed at the bottom of the table.

Table 6B: Nominal Bonds and Commodity Index Risk Premia

	FB 2y	FB 3y	FB 4y	FB 5y	CRB CI
Data	0.69%	1.25%	1.72%	2.10%	-1.45%
Model	0.56%	0.75%	1.02%	1.16%	-1.21%
Cash Flow	0.00%	-0.02%	0.02%	0.03%	0.01%
Inflation	0.33%	0.48%	0.59%	0.68%	-1.11%
Volatility	0.23%	0.29%	0.41%	0.45%	-0.12%
Model	λ_C	λ_{ECF}	$\lambda_{E\pi}$	λ_V	
Implied	γ	$\gamma - \frac{1}{\psi}$	$-(\gamma - \frac{1}{\psi})$	$-\frac{(\gamma-1)(\gamma-\frac{1}{\psi})}{(1-\frac{1}{\psi})^2}$	
Risk Price	4.13	3.67	-3.67	-38.61	

The table shows the model fit for 1-year holding period excess returns on the Fama-Bliss bonds as well as for a spot commodity index constructed by the Commodity Research Bureau. The first row is the excess return on the portfolio in the data. The second row shows the model implied risk premium for each portfolio. The following three rows give a breakdown of the model implied risk premium into 3 components (cash-flow risk, inflation risk and volatility risk) where cash-flow risk premium is defined as $\lambda_C \beta_{i,C} + \lambda_{ECF} \beta_{i,ECF}$, expected inflation risk premium as $\lambda_{E\pi} \beta_{i,E\pi}$ and volatility premium as $\lambda_V \beta_{i,V}$. The model implied market prices of risk are displayed at the bottom of the table.

Table 6C: Size and Book to Market Portfolios Risk Premia

	S1	S2	S3	S4	S5
Data	5.32%	6.96%	6.83%	8.00%	9.79%
Model	5.44%	5.91%	5.38%	6.24%	8.37%
Cash Flow	1.72%	1.75%	1.60%	1.27%	1.56%
Inflation	1.59%	1.33%	1.03%	1.53%	1.93%
Volatility	2.13%	2.83%	2.76%	3.45%	4.87%
	BM1	BM2	BM3	BM4	BM5
Data	7.95%	7.90%	7.99%	7.79%	5.79%
Model	7.83%	7.80%	7.76%	7.05%	6.00%
Cash Flow	2.00%	1.89%	1.73%	1.65%	1.92%
Inflation	1.77%	1.78%	1.85%	1.65%	1.46%
Volatility	4.06%	4.13%	4.17%	3.75%	2.62%
Model	λ_C	λ_{ECF}	$\lambda_{E\pi}$	λ_V	
Implied	γ	$\gamma - \frac{1}{\psi}$	$-(\gamma - \frac{1}{\psi})$	$-\frac{(\gamma-1)(\gamma-\frac{1}{\psi})}{(1-\frac{1}{\psi})^2}$	
Risk Price	4.13	3.67	-3.67	-38.61	

The table shows the model fit for the 5 quintile size portfolios as well as the 5 quintile book to market portfolios taken from Kenneth French's website. The first row is the excess return on the portfolio in the data. The second row shows the model implied risk premium for each portfolio. The following three rows give a breakdown of the model implied risk premium into 3 components (cash-flow risk, inflation risk and volatility risk) where cash-flow risk premium is defined as $\lambda_C \beta_{i,C} + \lambda_{ECF} \beta_{i,ECF}$, expected inflation risk premium as $\lambda_{E\pi} \beta_{i,E\pi}$ and volatility premium as $\lambda_V \beta_{i,V}$. The model implied market prices of risk are displayed at the bottom of the table.

7 Figures

Figure 1

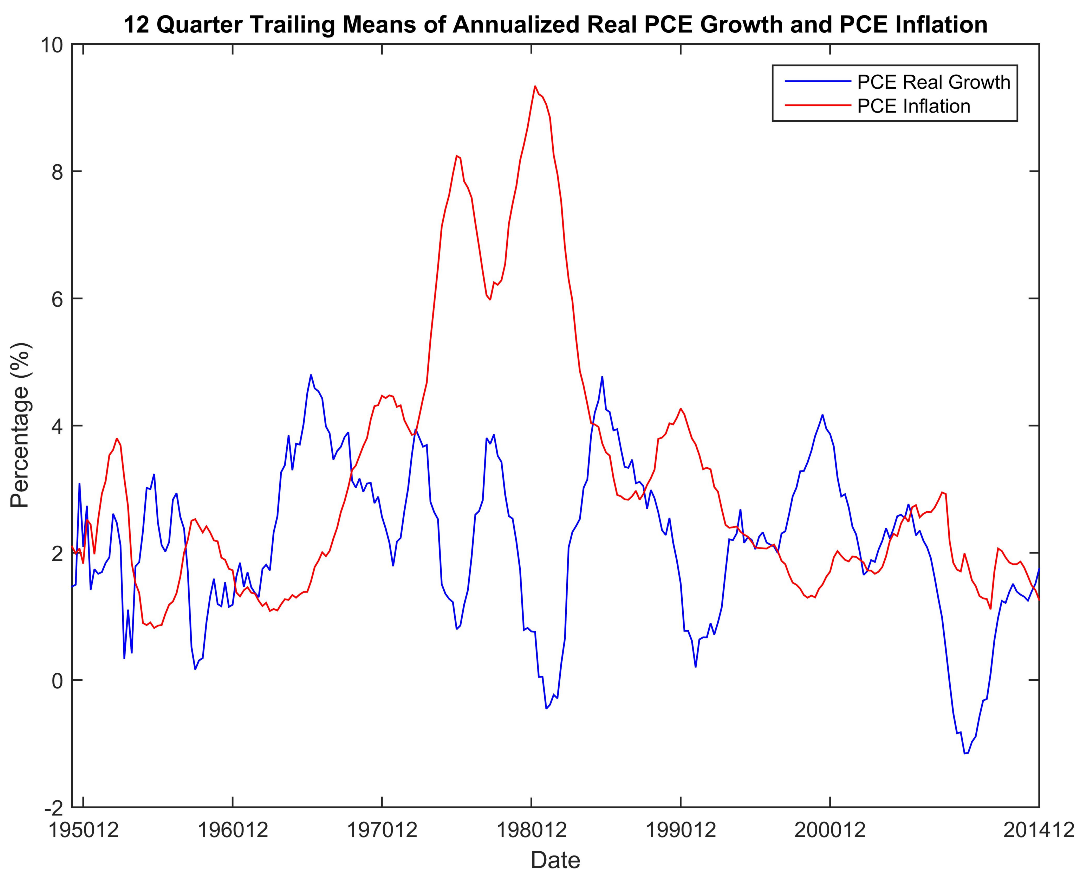

Figure 2

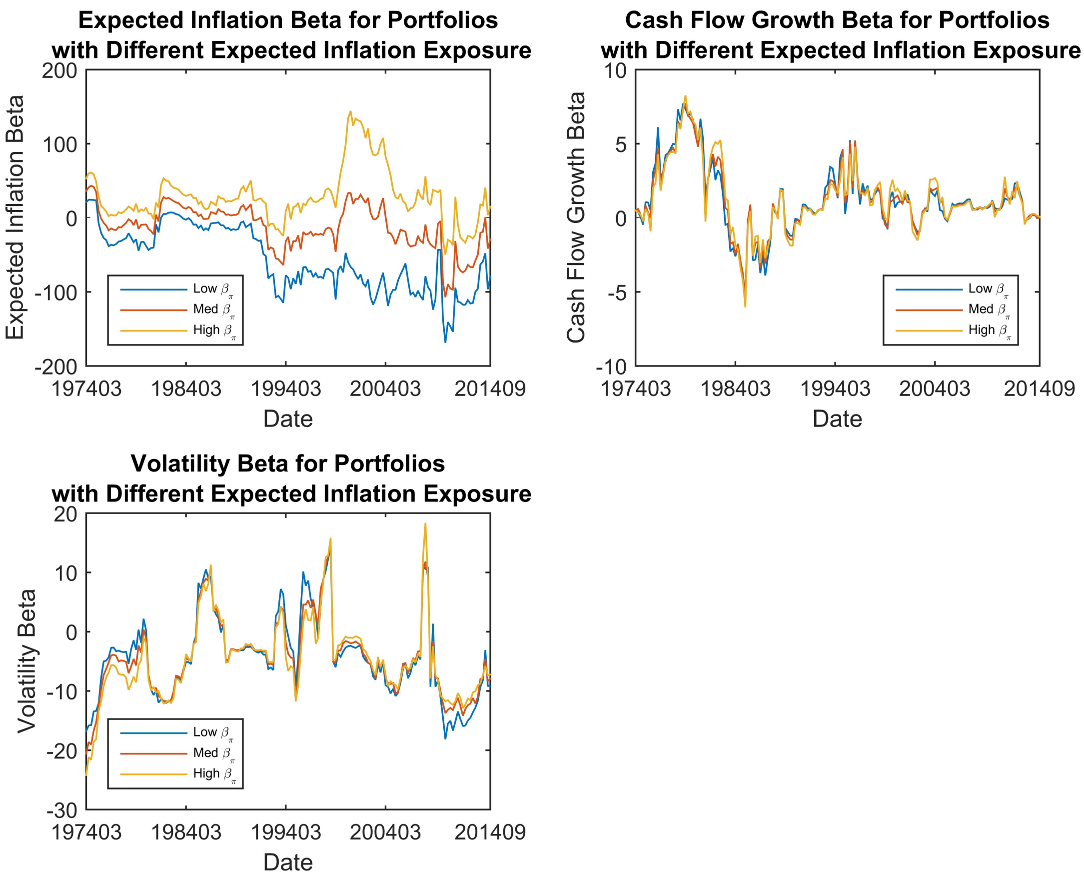
Expected Inflation Beta for Portfolios
with Different Expected Inflation Exposure
Expected Inflation Beta
200
100
0
-100
-200
197403
198403
199403
200403
201409
Date
Low β_π
Med β_π
High β_π
Cash Flow Growth Beta for Portfolios
with Different Expected Inflation Exposure
Cash Flow Growth Beta
10
5
0
-5
-10
197403
198403
199403
200403
201409
Date
Low β_π
Med β_π
High β_π
Volatility Beta for Portfolios
with Different Expected Inflation Exposure
Volatility Beta
20
10
0
-10
-20
-30
197403
198403
199403
200403
201409
Date
Low β_π
Med β_π
High β_π

Figure 3

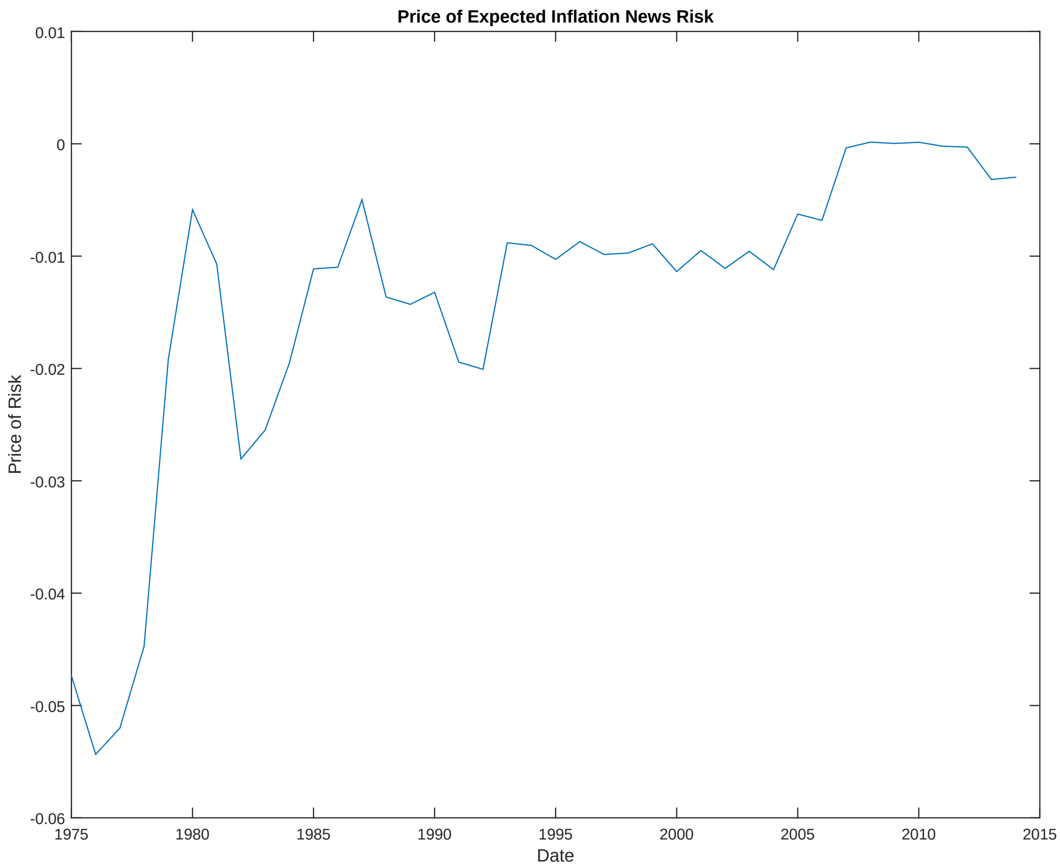
Price of Expected Inflation News Risk
0.01
0
-0.01
-0.02
-0.03
-0.04
-0.05
-0.06
Price of Risk
1975
1980
1985
1990
1995
2000
2005
2010
2015
Date